

Composed September 5, 2002



System Administration Guide: Resource Management and Network Services

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

Part No: 816-7125-10
December 2002

Composed September 5, 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, UltraSPARC, WebNFS, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

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

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

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



020905 @ 4660



Contents

Preface	35
1 Resource Management and Network Service (Overview)	39
Topics for the Solaris 9 Release	39
Perl 5	40
Accessing Perl Documentation	40
Perl Compatibility Issues	40
Changes to the Solaris Version of Perl	41
2 Managing Web Cache Servers	43
Network Cache and Accelerator (Overview)	43
New NCA Functionality	44
Managing Web Cache Servers (Task Map)	44
Planning for NCA	45
System Requirements for NCA	45
NCA Logging	45
Interpositioning Library for Daemon Support of the Door Server	45
Multiple Instance Support	46
Administering the Caching of Web Pages (Tasks)	46
▼ How to Enable Caching of Web Pages	46
▼ How to Disable Caching of Web Pages	48
▼ How to Enable or Disable NCA Logging	49
▼ How to Load the Socket Utility Library for NCA	49
▼ How to Add a New Port to the NCA Service	50
Caching Web Pages (Reference)	50

Composed September 5, 2002

	NCA Files	50
	NCA Architecture	52
3	Time-Related Services	55
	Clock Synchronization (Overview)	55
	Managing Network Time Protocol (Tasks)	56
	▼ How to Set Up an NTP Server	56
	▼ How to Set Up an NTP Client	56
	Using Other Time-Related Commands (Tasks)	57
	▼ How to Synchronize Date and Time From Another System	57
	Network Time Protocol (Reference)	58
4	Solaris 9 Resource Manager Topics	59
5	Introduction to Solaris 9 Resource Manager	61
	Overview	61
	Resource Classifications	62
	Resource Management Control Mechanisms	63
	Resource Management Configuration	64
	When to Use Resource Management	64
	Server Consolidation	64
	Supporting a Large or Varied User Population	65
	Setting Up Resource Management (Task Map)	66
6	Projects and Tasks	69
	Overview	69
	Projects	70
	Determining a User's Default Project	70
	project Database	70
	PAM Subsystem	71
	Name Service Configuration	71
	Local project File Format	71
	Name Service Configuration for NIS	73
	Directory Service Configuration for LDAP	73
	Tasks	74
	Commands Used to Administer Projects and Tasks	74

Composed September 5, 2002

	Command Options Used With Projects and Tasks	75
	Using <code>cron</code> and <code>su</code> With Projects and Tasks	77
	Project Administration Examples	77
	▼ How to Define a Project and View the Current Project	77
	▼ How to Delete a Project From the <code>/etc/project</code> File	78
	▼ How to Obtain User and Project Membership Information	79
	▼ How to Create a New Task	79
	▼ How to Move a Running Process Into a New Task	79
7	Extended Accounting	81
	Overview	81
	How Extended Accounting Works	82
	Extensible Format	83
	<code>exacct</code> Records and Format	83
	Extended Accounting Configuration	84
	Commands Used With Extended Accounting	84
	Using Extended Accounting Functionality	85
	▼ How to Activate Extended Accounting for Processes, Tasks, and Flows	85
	▼ How to Activate Extended Accounting With a Startup Script	85
	▼ How to Display Extended Accounting Status	85
	▼ How to View Available Accounting Resources	86
	▼ How to Deactivate Process, Task, and Flow Accounting	87
8	Resource Controls	89
	Overview	89
	Administering Resource Controls	90
	Available Resource Controls	90
	Resource Control Values and Privilege Levels	91
	Actions on Resource Control Values	92
	Resource Control Flags and Properties	93
	Resource Control Enforcement	94
	Global Monitoring of Resource Control Events	94
	Configuration	95
	Temporarily Updating Resource Control Values on a Running System	95
	Updating Logging Status	95
	Updating Resource Controls	95
	Using Resource Controls	96

▼ How to Set the Maximum Number of LWPs for Each Task in a Project	96
▼ How to Set Multiple Controls on a Project	97
▼ How to Use <code>prctl</code>	97
▼ How to Use <code>rctladm</code>	98
Capacity Warnings	98
▼ How to Determine Whether a Web Server Is Allocated Enough CPU Capacity	98
9 Fair Share Scheduler	101
Overview	101
CPU Share Definition	102
CPU Shares and Process State	102
CPU Share Versus Utilization	103
CPU Share Examples	103
Example 1: Two CPU-Bound Processes in Each Project	104
Example 2: No Competition Between Projects	104
Example 3: One Project Unable to Run	105
FSS Configuration	106
Projects and Users	106
CPU Shares Configuration	106
FSS and Processor Sets	107
FSS and Processor Sets Examples	108
Combining FSS With Other Scheduling Classes	110
Monitoring the FSS	111
▼ How to Monitor System CPU Usage by Projects	111
▼ How to Monitor CPU Usage by Projects in Processor Sets	111
FSS Configuration Examples	111
▼ How to Set the Scheduler Class	111
▼ How to Manually Move Processes From the TS Into the FSS Class	112
▼ How to Manually Move Processes From all User Classes Into the FSS Class	112
▼ How to Move a Project's Processes Into the FSS Class	112
▼ How to Tune Scheduler Parameters	113
References	114
10 Resource Pools	115
Overview	115

When to Use Pools	116
Batch Compute Server	116
Application or Database Server	116
Turning on Applications in Phases	116
Complex Timesharing Server	116
Workloads That Change Seasonally	117
Real-Time Applications	117
Administering Pools	117
Pools Framework	118
Implementing Pools on a System	118
Dynamic Reconfiguration Operations and Resource Pools	118
Creating Pools Configurations	119
▼ How to Create a Configuration by Discovery	120
▼ How to Create a New Configuration	120
▼ How to Modify a Configuration	120
▼ How to Associate a Pool With a Scheduling Class	121
▼ How to Use Command Files With <code>poolcfg</code>	122
Activating and Deactivating Pools Configurations	123
▼ How to Activate a Pools Configuration	123
▼ How to Deactivate a Pools Configuration	123
Binding to a Pool	124
▼ How to Bind Processes to a Pool	124
▼ How to Bind Tasks or Projects to a Pool	124
▼ How to Use <code>project</code> Attributes to Bind New Processes to a Pool	124
▼ How to Use <code>project</code> Attributes to Bind a Process to a Different Pool	125
11 Resource Management Configuration Example	127
Configuration to Be Consolidated	127
Consolidation Configuration	128
Creating the Configuration	128
Viewing the Configuration	129
12 Resource Control Functionality in the Solaris Management Console	133
Using the Console (Task Map)	133
Overview	134
Management Scope	134
Performance Tool	134

▼ How to Access the Performance Tool	135
Monitoring by System	136
Monitoring by Project or User Name	136
Resource Controls Tab	138
▼ How to Access the Resource Controls Tab	139
Resource Controls You Can Set	140
Setting Values	140
References	140
13 Accessing Remote File Systems Topics	141
14 Managing Network File Systems (Overview)	143
NFS Terminology	143
NFS Servers and Clients	143
NFS File Systems	144
About the NFS Service	144
About Autofs	145
Features of the NFS Service	146
NFS Version 2 Protocol	146
NFS Version 3 Protocol	146
NFS ACL Support	147
NFS Over TCP	147
Network Lock Manager and NFS	147
NFS Large File Support	147
NFS Client Failover	147
Kerberos Support for the NFS Service	148
WebNFS Support	148
RPCSEC_GSS Security Flavor	148
Solaris 7 Extensions for NFS Mounting	148
Security Negotiation for the WebNFS Service	149
NFS Server Logging	149
Autofs Features	149
15 Remote File-System Administration (Tasks)	151
Automatic File-System Sharing	152
▼ How to Set Up Automatic File-System Sharing	152
▼ How to Enable WebNFS Access	154

▼ How to Enable NFS Server Logging	155
Mounting File Systems	156
▼ How to Mount a File System at Boot Time	157
▼ How to Mount a File System From the Command Line	158
Mounting With the Automounter	159
▼ How to Disable Large Files on an NFS Server	159
▼ How to Use Client-Side Failover	160
▼ How to Disable Mount Access for One Client	160
▼ How to Mount an NFS File System Through a Firewall	161
▼ How to Mount an NFS File System Using an NFS URL	161
Setting Up NFS Services	162
▼ How to Start the NFS Services	162
▼ How to Stop the NFS Services	162
▼ How to Start the Automounter	163
▼ How to Stop the Automounter	163
Administering the Secure NFS System	163
▼ How to Set Up a Secure NFS Environment With DH Authentication	164
WebNFS Administration Tasks	165
Planning for WebNFS Access	166
▼ How to Browse Using an NFS URL	167
▼ How to Enable WebNFS Access Through a Firewall	167
Task Overview for Autofs Administration	168
Task Map for Autofs Administration	168
Administrative Tasks Involving Maps	170
Modifying the Maps	171
▼ How to Modify the Master Map	171
▼ How to Modify Indirect Maps	171
▼ How to Modify Direct Maps	171
Avoiding Mount-Point Conflicts	172
Accessing Non-NFS File Systems	172
How to Access CD-ROM Applications With Autofs	173
▼ How to Access PC-DOS Data Diskettes With Autofs	173
Accessing NFS File Systems Using CacheFS	173
▼ How to Access NFS File Systems Using CacheFS	174
Customizing the Automounter	174
▼ Setting Up a Common View of /home	174
▼ How to Set Up /home With Multiple Home Directory File Systems	175
▼ How to Consolidate Project-Related Files Under /ws	176

▼ How to Set Up Different Architectures to Access a Shared Namespace	177
▼ How to Support Incompatible Client Operating System Versions	178
▼ How to Replicate Shared Files Across Several Servers	179
▼ How to Apply Autofs Security Restrictions	179
▼ How to Use a Public File Handle With Autofs	179
▼ How to Use NFS URLs With Autofs	180
Disabling Autofs Browsability	180
▼ How to Completely Disable Autofs Browsability on a Single NFS Client	180
▼ How to Disable Autofs Browsability for All Clients	181
▼ How to Disable Autofs Browsability on a Selected File System	181
Strategies for NFS Troubleshooting	182
NFS Troubleshooting Procedures	183
▼ How to Check Connectivity on an NFS Client	183
▼ How to Check the NFS Server Remotely	184
▼ How to Verify the NFS Service on the Server	185
▼ How to Restart NFS Services	187
▼ How to Warm-Start rpcbind	187
▼ Identifying Which Host Is Providing NFS File Service	188
▼ How to Verify Options Used With the mount Command	188
Troubleshooting Autofs	189
Error Messages Generated by automount -v	189
Miscellaneous Error Messages	190
Other Errors With Autofs	191
NFS Error Messages	191
16 Accessing Remote File Systems (Reference)	197
NFS Files	197
/etc/default/nfslogd	198
/etc/nfs/nfslog.conf	199
NFS Daemons	200
automountd	200
lockd	201
mountd	202
nfsd	202
nfslogd	203
statd	203
NFS Commands	204

Composed September 5, 2002

automount	204
clear_locks	205
mount	205
umount	208
mountall	209
umountall	209
share	210
unshare	214
shareall	215
unshareall	215
showmount	215
setmnt	216
Other Useful Commands	217
nfsstat	217
pstack	218
rpcinfo	219
snoop	221
truss	221
How the NFS Service Works	222
Version 2 and Version 3 Negotiation	222
UDP and TCP Negotiation	222
File Transfer Size Negotiation	223
How File Systems Are Mounted	223
Effects of the -public Option and NFS URLs When Mounting	224
Client-Side Failover	224
Large Files	226
How NFS Server Logging Works	226
How the WebNFS Service Works	227
WebNFS Limitations With Web Browser Use	228
Secure NFS System	229
Secure RPC	229
Autofs Maps	232
Master Autofs Map	232
Direct Autofs Maps	234
Indirect Autofs Maps	236
How Autofs Works	238
How Autofs Navigates Through the Network (Maps)	239
How Autofs Starts the Navigation Process (Master Map)	240

	Autofs Mount Process	240
	How Autofs Selects the Nearest Read-Only Files for Clients (Multiple Locations)	242
	Variables in a Map Entry	244
	Maps That Refer to Other Maps	245
	Executable Autofs Maps	247
	Modifying How Autofs Navigates the Network (Modifying Maps)	247
	Default Autofs Behavior With Name Services	247
	Autofs Reference	249
	Metacharacters	249
	Special Characters	250
17	SLP Topics	251
18	SLP (Overview)	253
	SLP Architecture	253
	Summary of the SLP Design	254
	SLP Agents and Processes	254
	SLP Implementation	256
	Other SLP Information Sources	257
19	Planning and Enabling SLP (Tasks)	259
	SLP Configuration Considerations	259
	Deciding What to Reconfigure	260
	Using snoop to Monitor SLP Activity	260
	▼ How to Use snoop to Run SLP Traces	261
	Analyzing a snoop slp Trace	261
	Enabling SLP	263
20	Administering SLP (Tasks)	265
	Configuring SLP Properties	265
	SLP Configuration File: Basic Elements	266
	▼ How to Change Your SLP Configuration	267
	Modifying DA Advertising and Discovery Frequency	268
	Limiting UAs and SAs to Statically Configured DAs	268
	▼ How to Limit UAs and SAs to Statically Configured DAs	269

Configuring DA Discovery for Dial-up Networks	269
▼ How to Configure DA Discovery for Dial-up Networks	270
Configuring the DA Heartbeat for Frequent Partitions	271
▼ How to Configure DA Heartbeat for Frequent Partitions	271
Relieving Network Congestion	272
Accommodating Different Network Media, Topologies, or Configurations	272
Reducing SA Reregistrations	273
▼ How to Reduce SA Reregistrations	273
Configuring the Multicast Time-to-Live Property	274
▼ How to Configure the Multicast Time-to-Live Property	274
Configuring the Packet Size	275
▼ How to Configure the Packet Size	275
Configuring Broadcast-Only Routing	276
▼ How to Configure Broadcast-Only Routing	276
Modifying Timeouts on SLP Discovery Requests	277
Changing Default Timeouts	277
▼ How to Change Default Timeouts	278
Configuring the Random-Wait Bound	279
▼ How to Configure the Random-Wait Bound	280
Deploying Scopes	281
When to Configure Scopes	281
Considerations When Configuring Scopes	282
▼ How to Configure Scopes	282
Deploying DAs	284
Why Deploy a SLP DA?	284
When to Deploy DAs	285
▼ How to Deploy DAs	285
Where to Place DAs	286
Multihoming	287
Multihoming Configuration	287
When to Configure for Nonrouted, Multiple Network Interfaces	287
Configuring Nonrouted, Multiple Network Interfaces (Task Map)	288
Configuring the <code>net.slp.interfaces</code> Property	288
Proxy Advertising on Multihomed Hosts	290
DA Placement and Scope Name Assignment	291
Considerations When Configuring for Nonrouted, Multiple Network Interfaces	291

21	Incorporating Legacy Services	293
	When to Advertise Legacy Services	293
	Advertising Legacy Services	293
	Modifying the Service	294
	Advertising a Service That Is Not SLP Enabled	294
	SLP Proxy Registration	294
	▼ How to Enable SLP Proxy Registration	294
	Using SLP Proxy Registration to Advertise	295
	Considerations When Advertising Legacy Services	297
22	SLP (Reference)	299
	SLP Status Codes	299
	SLP Message Types	301
23	Mail Services Topics	303
24	Mail Services (Overview)	305
	What's New in Version 8.12 of <code>sendmail</code>	305
	Other <code>sendmail</code> Information Sources	306
	Introduction to the Components of Mail Services	306
	Overview of the Software Components	307
	Overview of the Hardware Components	307
25	Mail Services (Tasks)	309
	Mail Services Task Map	310
	Planning Your Mail System	311
	Local Mail Only	312
	Local Mail and a Remote Connection	312
	Setting Up Mail Services (Task Map)	314
	Setting Up Mail Services (Tasks)	314
	▼ How to Set Up a Mail Server	315
	▼ How to Set Up a Mail Client	317
	▼ How to Set Up a Mail Host	319
	▼ How to Set Up a Mail Gateway	320
	▼ How to Use DNS With <code>sendmail</code>	322
	▼ How to Set Up a Virtual Host	322

Building the <code>sendmail.cf</code> Configuration File (Task)	323
▼ How to Build a New <code>sendmail.cf</code> File	323
Managing Mail Delivery by Using an Alternate Configuration (Task)	325
▼ How to Manage Mail Delivery by Using an Alternate Configuration of <code>sendmail.cf</code>	325
Administering Mail Alias Files (Task Map)	326
Administering Mail Alias Files (Tasks)	327
▼ How to Manage Alias Entries in an NIS+ <code>mail_aliases</code> Table	327
▼ How to Set Up an NIS <code>mail_aliases</code> Map	332
▼ How to Set Up a Local Mail Alias File	333
▼ How to Create a Keyed Map File	335
Managing the <code>postmaster</code> Alias	335
Administering the Queue Directories (Task Map)	338
Administering the Queue Directories (Tasks)	338
▼ How to Display the Contents of the Mail Queue, <code>/var/spool/mqueue</code>	339
▼ How to Force Mail Queue Processing in the Mail Queue, <code>/var/spool/mqueue</code>	339
▼ How to Run a Subset of the Mail Queue, <code>/var/spool/mqueue</code>	339
▼ How to Move the Mail Queue, <code>/var/spool/mqueue</code>	340
▼ How to Run the Old Mail Queue, <code>/var/spool/omqueue</code>	341
Administering <code>.forward</code> Files (Task Map)	341
Administering <code>.forward</code> Files (Tasks)	342
▼ How to Disable <code>.forward</code> Files	342
▼ How to Change the <code>.forward</code> File Search Path	342
▼ How to Create and Populate <code>/etc/shells</code>	343
Troubleshooting Procedures and Tips for Mail Services (Task Map)	344
Troubleshooting Procedures and Tips for Mail Services (Tasks)	344
▼ How to Test the Mail Configuration	344
▼ How to Check Mail Aliases	345
▼ How to Test the <code>sendmail</code> Rule Sets	346
▼ How to Verify Connections to Other Systems	347
Logging Error Messages	347
Other Sources for Mail Diagnostic Information	348
Resolving Error Messages	348
26 Mail Services (Reference)	351
Solaris Version of <code>sendmail</code>	351
Flags Used and Not Used to Compile <code>sendmail</code>	352

Composed September 5, 2002

Alternative sendmail Commands	353
Versions of the Configuration File	354
Software and Hardware Components of Mail Services	355
Software Components	355
Hardware Components	362
Mail Service Programs and Files	365
Contents of the /usr/bin Directory	366
Contents of the /etc/mail Directory	366
Contents of the /usr/lib Directory	367
Other Files Used for Mail Services	370
Interactions of Mail Programs	370
sendmail Program	371
Mail Alias Files	376
.forward Files	379
/etc/default/sendmail File	381
Mail Addresses and Mail Routing	382
Interactions of sendmail With Name Services	383
sendmail.cf and Mail Domains	383
sendmail and Name Services	384
Interactions of NIS and sendmail	385
Interactions of sendmail With NIS and DNS	386
Interactions of NIS+ and sendmail	386
Interactions of sendmail With NIS+ and DNS	387
27 What's New With Mail Services (Reference)	389
Changes to sendmail	389
New Configuration File, submit.cf	390
New or Deprecated Command-Line Options	391
New and Revised Configuration File Options and Related Topics	392
New Defined Macros for sendmail	406
New Macros Used to Build the sendmail Configuration File	407
New and Revised m4 Configuration Macros for sendmail	408
Changes to the FEATURE() Declaration	409
Changes to the MAILER() Declaration	412
New Delivery Agent Flags	412
New Equates for Delivery Agents	413
New Queue Features	414

Composed September 5, 2002

	Changes for LDAP in sendmail	415
	New Built-in Mailer Feature	416
	New Rule Sets	416
	Changes to Files	417
	IPv6 Addresses in Configuration	418
	Changes to mail.local	418
	Changes to mailstats	419
	Changes to makemap	419
	New Command, editmap	420
	Other Changes and Features of Interest	421
28	Modem-Related Network Services Topics	423
29	Solaris PPP 4.0 (Overview)	425
	Solaris PPP 4.0 Basics	425
	Solaris PPP 4.0 Compatibility	426
	Which Version of Solaris PPP to Use	426
	Where to Go for More Information About PPP	427
	PPP Configurations and Terminology	428
	Dial-up PPP Overview	429
	Leased-Line PPP Overview	432
	PPP Authentication	434
	Authenticators and Authenticatees	435
	PPP Authentication Protocols	435
	Why Use PPP Authentication?	436
	Support for DSL Users Through PPPoE	436
	PPPoE Overview	437
	Parts of a PPPoE Configuration	437
	Security on a PPPoE Tunnel	439
30	Planning for the PPP Link (Tasks)	441
	Overall PPP Planning (Task Map)	441
	Planning a Dial-up PPP Link	442
	Before You Set Up the Dial-out Machine	442
	Before You Set Up the Dial-in Server	443
	Example—Configuration for Dial-up PPP	443
	Where to Go For More Information About Dial-up PPP	446

Planning a Leased-Line Link	446
Before You Set Up the Leased-Line Link	446
Example—Configuration for a Leased-Line Link	447
Where to Get More Information About Leased Lines	449
Planning for Authentication on a Link	449
Before You Set Up PPP Authentication	449
Example—PPP Authentication Configurations	450
Where to Get More Information About Authentication	453
Planning for DSL Support Over a PPPoE Tunnel	454
Before You Set Up a PPPoE Tunnel	454
Example—Configuration for a PPPoE Tunnel	456
Where to Get More Information About PPPoE	457
31 Setting Up a Dial-up PPP Link (Tasks)	459
Major Tasks for Setting Up the Dial-up PPP Link (Task Map)	459
Configuring the Dial-out Machine	460
Tasks for Configuring the Dial-out Machine (Task Map)	460
Dial-up PPP Template Files	461
Configuring Devices on the Dial-out Machine	461
▼ How to Configure the Modem and Serial Port (Dial-out Machine)	462
Configuring Communications on the Dial-out Machine	463
▼ How to Define Communications Over the Serial Line	463
▼ How to Create the Instructions for Calling a Peer	464
▼ How to Define the Connection With an Individual Peer	465
Configuring the Dial-in Server	467
Tasks for Configuring the Dial-in Server (Task Map)	467
Configuring Devices on the Dial-in Server	468
How to Configure the Modem and Serial Port (Dial-in Server)	468
▼ How to Set the Modem Speed	469
Setting Up Users of the Dial-in Server	469
▼ How to Configure Users of the Dial-in Server	470
Configuring Communications Over the Dial-in Server	471
How to Define Communications Over the Serial Line (Dial-in Server)	471
Calling the Dial-in Server	473
▼ How to Call the Dial-in Server	473
Where to Go From Here	474

32	Setting Up a Leased-Line PPP Link (Tasks)	475
	Setting Up a Leased Line (Task Map)	475
	Configuring Synchronous Devices on the Leased Line	476
	Prerequisites for Synchronous Devices Setup	476
	▼ How to Configure Synchronous Devices	476
	Configuring a Machine on the Leased Line	477
	Prerequisites for Configuring the Local Machine on a Leased Line	478
	▼ How to Configure a Machine on a Leased Line	478
33	Setting Up PPP Authentication (Tasks)	481
	Configuring PPP Authentication (Task Map)	481
	Configuring PAP Authentication	482
	Setting Up PAP Authentication (Task Maps)	482
	Configuring PAP Authentication on the Dial-in Server	483
	▼ How to Create a PAP Credentials Database (Dial-in Server)	483
	Modifying the PPP Configuration Files for PAP (Dial-in Server)	485
	▼ How to Add PAP Support to the PPP Configuration Files (Dial-in Server)	485
	Configuring PAP Authentication for Trusted Callers (Dial-out Machines)	486
	▼ How to Configure PAP Authentication Credentials for the Trusted Callers	487
	Modifying PPP Configuration Files for PAP (Dial-out Machine)	488
	▼ How to Add PAP Support to the PPP Configuration Files (Dial-out Machine)	488
	Configuring CHAP Authentication	490
	Setting Up CHAP Authentication (Task Maps)	490
	Configuring CHAP Authentication on the Dial-in Server	491
	▼ How to Create a CHAP Credentials Database (Dial-in Server)	491
	Modifying the PPP Configuration Files for CHAP (Dial-in Server)	492
	▼ How to Add CHAP Support to the PPP Configuration Files (Dial-in Server)	492
	Configuring CHAP Authentication for Trusted Callers (Dial-out Machines)	493
	▼ How to Configure CHAP Authentication Credentials for the Trusted Callers	493
	▼ Adding CHAP to the Configuration Files (Dial-out Machine)	494
	How to Add CHAP Support to the PPP Configuration Files (Dial-out Machine)	494

34	Setting Up a PPPoE Tunnel (Tasks)	497
	Major Tasks for Setting Up a PPPoE Tunnel (Task Maps)	497
	Setting Up the PPPoE Client	498
	Prerequisites for Setting Up the PPPoE Client	498
	▼ How to Configure an Interface for a PPPoE Client	499
	▼ How to Define a PPPoE Access Server Peer	499
	Setting Up a PPPoE Access Server	501
	▼ How to Configure the Access Server's Interfaces for PPPoE	501
	▼ How to Provide Services to Clients of the Access Server	502
	▼ How to Modify an Existing <code>/etc/ppp/pppoe</code> File	502
	▼ How to Restrict the Use of an Interface to Particular Clients	502
	Where to Go From Here	504
35	Fixing Common Problems (Tasks)	505
	Solving PPP Problems (Task Map)	505
	Tools for Troubleshooting PPP	506
	▼ How to Obtain Diagnostic Information From <code>pppd</code>	507
	▼ How to Turn on PPP Debugging	508
	Fixing Network Problems That Affect PPP Performance	509
	▼ How to Diagnose Network Problems	509
	Fixing General Communications Problems	511
	▼ How to Diagnose and Fix Communications Problems	511
	Fixing PPP Configuration Problems	512
	▼ How to Diagnose Problems With the PPP Configuration	513
	Fixing Modem-Related Problems	513
	▼ How to Diagnose Modem Problems	514
	Fixing Chat Script-Related Problems	514
	▼ How to Obtain Debugging Information for Chat Scripts	515
	Fixing Serial-Line Speed Problems	517
	▼ How to Diagnose and Fix Serial-Line Speed Problems	517
	Fixing Leased-Line Problems	518
	Diagnosing and Fixing Authentication Problems	519
	Diagnosing and Fixing PPPoE Problems	519
	How to Obtain Diagnostic Information for PPPoE	519
36	Solaris PPP 4.0 Reference	523
	Using PPP Options in Files and on the Command Line	523

Where to Define PPP Options	523
How PPP Options Are Processed	525
How PPP Configuration File Privileges Work	525
/etc/ppp/options Configuration File	528
/etc/ppp/options.ttyname Configuration File	529
Configuring User-Specific Options	531
Configuring \$HOME/.ppprc on a Dial-in Server	531
Configuring \$HOME/.ppprc on a Dial-out Machine	532
Specifying Information for Communicating With the Dial-in Server	532
/etc/ppp/peers/peer-name File	532
/etc/ppp/peers/myisp.tmpl Template File	534
Where to Find Sample /etc/ppp/peers/peer-name Files	535
Configuring Modems for a Dial-up Link	535
Configuring the Modem Speed	535
Defining the Conversation on the Dial-up Link	536
Contents of the Chat Script	536
Chat Script Examples	537
Invoking the Chat Script	543
▼ How to Invoke a Chat Script (Task)	544
Creating a Chat File That Is Executable	545
▼ How to Create an Executable Chat Program	545
Authenticating Callers on a Link	545
Password Authentication Protocol (PAP)	546
Challenge-Handshake Authentication Protocol (CHAP)	549
Creating an IP Addressing Scheme for Callers	552
Assigning Dynamic IP Addresses to Callers	552
Assigning Static IP Addresses to Callers	553
Assigning IP Addresses by sPPP Unit Number	554
Creating PPPoE Tunnels for DSL Support	554
Files for Configuring Interfaces for PPPoE	555
PPPoE Access Server Commands and Files	557
PPPoE Client Commands and Files	562
37 Migrating From Asynchronous Solaris PPP to Solaris PPP 4.0 (Tasks)	565
Before Converting asppp Files	565
Example—/etc/asppp.cf Configuration File	566
Example—/etc/uucp/Systems File	566

Example—/etc/uucp/Devices File	567
Example—/etc/uucp/Dialers File	567
Running the asppp2pppd Conversion Script (Tasks)	568
Prerequisites	568
▼ How to Convert From asppp to Solaris PPP 4.0	568
▼ How to View the Results of the Conversion	569
38 Overview of UUCP	571
UUCP Hardware Configurations	571
UUCP Software	572
UUCP Daemons	572
UUCP Administrative Programs	573
UUCP User Programs	574
UUCP Database Files	574
Configuring UUCP Database Files	575
39 Administering UUCP	577
UUCP Administration Task Map	577
Adding UUCP Logins	578
▼ How to Add UUCP Logins	578
Starting UUCP	579
▼ How to Start UUCP	579
uudemon.poll Shell Script	580
uudemon.hour Shell Script	580
uudemon.admin Shell Script	580
uudemon.cleanup Shell Script	580
Running UUCP Over TCP/IP	581
▼ How to Activate UUCP for TCP/IP	581
UUCP Security and Maintenance	582
Setting Up UUCP Security	582
Regular UUCP Maintenance	582
Troubleshooting UUCP	583
▼ How to Check for Faulty Modems or ACUs	583
▼ How to Debug Transmissions	583
Checking the UUCP /etc/uucp/Systems File	584
Checking UUCP Error Messages	585
Checking Basic Information	585

40	UUCP Reference	587
	UUCP /etc/uucp/Systems File	587
	UUCP System-Name Field	588
	UUCP Time Field	588
	UUCP Type Field	589
	UUCP Speed Field	589
	UUCP Phone Field	590
	UUCP Chat Script Field	590
	UUCP Hardware Flow Control	593
	UUCP Setting Parity	593
	UUCP /etc/uucp/Devices File	593
	UUCP Type Field	594
	UUCP Line Field	595
	UUCP Line2 Field	595
	UUCP Class Field	595
	UUCP Dialer-Token-Pairs Field	596
	UUCP Protocol Definitions in the Devices File	598
	UUCP /etc/uucp/Dialers File	599
	UUCP Hardware Flow Control	602
	UUCP Setting Parity	602
	Other Basic UUCP Configuration Files	603
	UUCP /etc/uucp/Dialcodes File	603
	UUCP /etc/uucp/Sysfiles File	604
	UUCP /etc/uucp/Sysname File	605
	UUCP /etc/uucp/Permissions File	605
	UUCP Structuring Entries	605
	UUCP Considerations	606
	UUCP REQUEST Option	606
	UUCP SENDFILES Option	607
	UUCP MYNAME Option	607
	UUCP READ and WRITE Options	608
	UUCP NOREAD and NOWRITE Options	608
	UUCP CALLBACK Option	608
	UUCP COMMANDS Option	609
	UUCP VALIDATE Option	610
	UUCP MACHINE Entry for OTHER	612
	Combining MACHINE and LOGNAME Entries for UUCP	612
	UUCP Forwarding	613

UUCP /etc/uucp/Poll File	613
UUCP /etc/uucp/Config File	613
UUCP/etc/uucp/Grades File	614
UUCP User-job-grade Field	614
UUCP System-job-grade Field	614
UUCP Job-size Field	615
UUCP Permit-type Field	615
UUCP ID-list Field	616
Other UUCP Configuration Files	616
UUCP /etc/uucp/Devconfig File	616
UUCP /etc/uucp/Limits File	617
UUCP remote.unknown File	617
UUCP Administrative Files	618
UUCP Error Messages	619
UUCP ASSERT Error Messages	619
UUCP STATUS Error Messages	621
UUCP Numerical Error Messages	622
41 Working With Remote Systems Topics	625
42 Working With Remote Systems (Overview)	627
What Is the FTP Server?	627
What Is a Remote System?	627
What's New for the Solaris 9 FTP Server?	628
43 Administering the FTP Server (Tasks)	631
Controlling FTP Server Access	632
▼ How to Define FTP Server Classes	633
▼ How to Set User Login Limits	634
▼ How to Control the Number of Invalid Login Attempts	635
▼ How to Disallow FTP Server Access to Particular Users	636
▼ How to Restrict Access to the Default FTP Server	637
Setting Up FTP Server Logins	638
▼ How to Set Up Real FTP Users	638
▼ How to Set Up Guest FTP Users	639
▼ How to Set Up Anonymous FTP Users	640
▼ How to Create the /etc/shells file	640

Customizing Message Files	641
▼ How to Customize Message Files	641
▼ How to Create Messages to Be Sent to Users	642
▼ How to Configure the README Option	643
Controlling Access to Files on the FTP Server	644
▼ How to Control File Access Commands	645
Controlling Uploads and Downloads on the FTP Server	645
▼ How to Control Uploads to the FTP Server	646
▼ How to Control Downloads to the FTP Server	648
Virtual Hosting	648
▼ How to Enable Limited Virtual Hosting	649
▼ How to Enable Complete Virtual Hosting	650
Starting the FTP Server Automatically	652
Starting an FTP Server From <code>inetd.conf</code>	652
▼ How to Start an FTP Server From <code>inetd.conf</code>	652
Starting a Standalone FTP Server	653
▼ How to Start a Standalone FTP Server	653
Shutting Down the FTP Server	653
▼ How to Shut Down the FTP Server	653
Debugging the FTP Server	654
▼ How to Check <code>syslogd</code> for FTP Server Messages	655
▼ How to Use <code>greeting text</code> to Verify <code>ftppass</code>	655
▼ How to Check the Commands Executed by FTP Users	655
44 Accessing Remote Systems (Tasks)	657
Logging In to a Remote System (<code>rlogin</code>)	658
Authentication for Remote Logins (<code>rlogin</code>)	658
Linking Remote Logins	661
Direct or Indirect Remote Logins	661
What Happens After You Log In Remotely	662
▼ How to Search for and Remove <code>.rhosts</code> Files	662
▼ How to Find Out If a Remote System Is Operating	663
▼ How to Find Who Is Logged In to a Remote System	664
▼ How to Log In to a Remote System (<code>rlogin</code>)	665
▼ How to Log Out From a Remote System (<code>exit</code>)	665
Logging In to a Remote System (<code>ftp</code>)	666
Authentication for Remote Logins (<code>ftp</code>)	666

Composed September 5, 2002

Essential ftp Commands	666
▼ How to Open an ftp Connection to a Remote System	667
▼ How to Close an ftp Connection to a Remote System	668
▼ How to Copy Files From a Remote System (ftp)	668
▼ How to Copy Files to a Remote System (ftp)	670
Remote Copying With rcp	673
Security Considerations for Copy Operations	673
Specifying Source and Target	673
▼ How to Copy Files Between a Local and a Remote System (rcp)	675
45 Monitoring Network Services Topics	679
46 Monitoring Network Performance (Tasks)	681
Monitoring Network Performance	681
▼ How to Check the Response of Hosts on the Network	682
▼ How to Send Packets to Hosts on the Network	683
▼ How to Capture Packets From the Network	683
▼ How to Check the Network Status	683
▼ How to Display NFS Server and Client Statistics	686
A System Administration Guide: Resource Management and Network Services Updates	691
Solaris 9 9/02 Updates	691
Solaris 9 12/02 Updates	691
Glossary	693
Index	699

Tables

TABLE 2-1	NCA Files	50
TABLE 3-1	NTP Files	58
TABLE 8-1	Standard Resource Controls	91
TABLE 8-2	Signals Available to Resource Control Values	92
TABLE 15-1	File-System Sharing Task Map	152
TABLE 15-2	Task Map for Mounting File Systems	157
TABLE 15-3	Task Map for NFS Services	162
TABLE 15-4	Task Map for WebNFS Administration	166
TABLE 15-5	Task Map for Autofs Administration	168
TABLE 15-6	Types of autofs Maps and Their Uses	170
TABLE 15-7	Map Maintenance	170
TABLE 15-8	When to Run the automount Command	170
TABLE 16-1	NFS Files	197
TABLE 16-2	NFS Security Modes	207
TABLE 16-3	Predefined Map Variables	245
TABLE 18-1	SLP Agents	254
TABLE 20-1	SLP Configuration Operations	266
TABLE 20-2	DA Advertisement Timing and Discovery Request Properties	268
TABLE 20-3	SLP Performance Properties	272
TABLE 20-4	Time-out Properties	277
TABLE 20-5	Configuring Nonrouted, Multiple Network Interfaces	288
TABLE 21-1	SLP Proxy Registration File Description	296
TABLE 22-1	SLP Status Codes	299
TABLE 22-2	SLP Message Types	301
TABLE 26-1	General sendmail Flags	352
TABLE 26-2	Maps and Database Types	352

TABLE 26-3	Solaris Flags	352
TABLE 26-4	Generic Flags Not Used in the Solaris Version of sendmail	353
TABLE 26-5	Alternate sendmail Commands	354
TABLE 26-6	Configuration File Version Values	354
TABLE 26-7	Top-Level Domains	358
TABLE 26-8	Conventions for the Format of Mailbox Names	360
TABLE 26-9	Contents of the /usr/lib Directory	367
TABLE 26-10	Contents of the /usr/lib/mail Directory Used for Mail Services	368
TABLE 26-11	Other Files Used for Mail Services	370
TABLE 26-12	Columns in the NIS+ mail_aliases Table	379
TABLE 27-1	New Command-Line Options for sendmail	392
TABLE 27-2	New and Revised Options for sendmail	393
TABLE 27-3	Deprecated and Unsupported Configuration File Options for sendmail	400
TABLE 27-4	New Keys for ClientPortOptions	401
TABLE 27-5	New and Revised Keys for DaemonPortOptions	402
TABLE 27-6	Values for the New Modifier Key	402
TABLE 27-7	Arguments for the PidFile and ProcessTitlePrefix Options	403
TABLE 27-8	New and Revised Arguments for PrivacyOptions	404
TABLE 27-9	New and Revised Settings for Timeout	405
TABLE 27-10	Defined Macros for sendmail	406
TABLE 27-11	New Macros Used to Build the sendmail Configuration File	407
TABLE 27-12	New MAX Macros	408
TABLE 27-13	New and Revised m4 Configuration Macros for sendmail	408
TABLE 27-14	New and Revised FEATURE() Declarations	409
TABLE 27-15	Unsupported FEATURE() Declarations	412
TABLE 27-16	New Mailer Flags	413
TABLE 27-17	New Equates for Delivery Agents	414
TABLE 27-18	Comparison of Tokens	415
TABLE 27-19	New LDAP Map Flags	416
TABLE 27-20	Possible Values for the First Mailer Argument	416
TABLE 27-21	New Rule Sets	416
TABLE 27-22	New Command-Line Options for mail.local	418
TABLE 27-23	New mailstats Options	419
TABLE 27-24	New makemap Options	419
TABLE 30-1	Task Map for PPP Planning	441
TABLE 30-2	Information for a Dial-out Machine	443

TABLE 30-3	Information for a Dial-in Server	443
TABLE 30-4	Planning for a Leased-Line Link	447
TABLE 30-5	Prerequisites Before Configuring Authentication	450
TABLE 30-6	Planning for PPPoE Clients	455
TABLE 30-7	Planning for a PPPoE Access Server	455
TABLE 31-1	Task Map for Setting Up the Dial-up PPP Link	459
TABLE 31-2	Task Map for Setting Up the Dial-out Machine	460
TABLE 31-3	Modem Settings for Dial-up PPP	462
TABLE 31-4	Task Map for Setting Up the Dial-in Server	467
TABLE 32-1	Task Map for Setting Up the Leased-Line Link	475
TABLE 33-1	Task Map for General PPP Authentication	482
TABLE 33-2	Task Map for PAP Authentication (Dial-in Server)	482
TABLE 33-3	Task Map for PAP Authentication (Dial-out Machine)	483
TABLE 33-4	Task Map for CHAP Authentication (Dial-in Server)	490
TABLE 33-5	Task Map for CHAP Authentication (Dial-out Machine)	490
TABLE 34-1	Task Map for Setting Up a PPPoE Client	497
TABLE 34-2	Task Map for Setting Up a PPPoE Access Server	498
TABLE 35-1	Task Map for Troubleshooting PPP	506
TABLE 35-2	Common Network Problems That Affect PPP	510
TABLE 35-3	General Communications Problems That Affect PPP	512
TABLE 35-4	Common PPP Configuration Problems	513
TABLE 35-5	Common Chat Script Problems	515
TABLE 35-6	Common Leased-Line Problems	518
TABLE 35-7	General Authentication Problems	519
TABLE 36-1	Summary of PPP Configuration Files and Commands	524
TABLE 36-2	Examples of the <code>/etc/ppp/options</code> File	529
TABLE 36-3	Examples of the <code>/etc/ppp/options.ttyname</code> File	531
TABLE 36-4	Examples of <code>/etc/ppp/peers/peer-name</code> Files	535
TABLE 36-5	Syntax of <code>/etc/ppp/pap-secrets</code>	546
TABLE 36-6	<code>/etc/ppp/pap-secrets</code> With login Option	549
TABLE 36-7	Syntax of <code>/etc/ppp/chap-secrets</code>	550
TABLE 36-8	PPPoE Commands and Configuration Files	555
TABLE 39-1	Task Map: UUCP Administration	577
TABLE 40-1	Day Field	588
TABLE 40-2	Escape Characters Used in Systems File Chat Script	591
TABLE 40-3	Dialer-Token Pairs	596
TABLE 40-4	Protocols Used in <code>/etc/uucp/Devices</code>	598
TABLE 40-5	Backslash Characters for <code>/etc/uucp/Dialers</code>	601

TABLE 40-6	Correspondences Between Dialcodes and Systems Files	603
TABLE 40-7	Entries in the Dialcodes File	603
TABLE 40-8	Job-size Field	615
TABLE 40-9	Permit-type Field	616
TABLE 40-10	UUCP Lock Files	618
TABLE 40-11	ASSERT Error Messages	620
TABLE 40-12	UUCP STATUS Messages	621
TABLE 40-13	UUCP Error Messages by Number	622
TABLE 42-1	What's New for the Solaris 9 FTP Server	628
TABLE 43-1	Task Map: Administering the FTP Server	631
TABLE 44-1	Task Map: Accessing Remote Systems	657
TABLE 44-2	Dependencies Between Login Method and Authentication Method (rlogin)	661
TABLE 44-3	Essential ftp Commands	666
TABLE 44-4	Allowed Syntaxes for Directory and File Names	674
TABLE 46-1	Network Monitoring Commands	681
TABLE 46-2	Output From the netstat -r Command	686
TABLE 46-3	Commands for Displaying Client/Server Statistics	687
TABLE 46-4	Output From the nfsstat -c Command	688
TABLE 46-5	Output From the nfsstat -m Command	689

Figures

FIGURE 2-1	Data Flow With the NCA Service	52
FIGURE 6-1	Project and Task Tree	74
FIGURE 7-1	Task Tracking With Extended Accounting Activated	82
FIGURE 8-1	Process Collectives, Container Relationships, and Their Resource Control Sets	94
FIGURE 9-1	FSS Scheduler Share Calculation	102
FIGURE 9-2	FSS Scheduler Share Calculation With Processor Sets	107
FIGURE 10-1	Resource Pool Framework	115
FIGURE 11-1	Server Consolidation Configuration	131
FIGURE 12-1	Performance Tool in the Solaris Management Console	134
FIGURE 12-2	Resource Controls Tab in the Solaris Management Console	138
FIGURE 16-1	/etc/init.d/autofs Script Starts automount	239
FIGURE 16-2	Navigation Through the Master Map	240
FIGURE 16-3	Server Proximity	243
FIGURE 16-4	How Autofs Uses the Name Service	248
FIGURE 18-1	SLP Basic Agents and Processes	254
FIGURE 18-2	SLP Architectural Agents and Processes Implemented With a DA	255
FIGURE 18-3	SLP Implementation	256
FIGURE 24-1	Typical Electronic Mail Configuration	307
FIGURE 25-1	Local Mail Configuration	312
FIGURE 25-2	Local Mail Configuration With a UUCP Connection	312
FIGURE 26-1	Gateway Between Different Communications Protocols	364
FIGURE 26-2	Interactions of Mail Programs	370
FIGURE 26-3	How sendmail Uses Aliases	373
FIGURE 26-4	Interaction of sendmail With Other Mail Programs	375
FIGURE 29-1	Parts of the PPP Link	428

Composed September 5, 2002

FIGURE 29-2	Basic Analog Dial-up PPP Link	430
FIGURE 29-3	Basic Leased-Line Configuration	432
FIGURE 29-4	Participants in a PPPoE Tunnel	437
FIGURE 30-1	Sample Dial-up Link	444
FIGURE 30-2	Sample Leased-Line Configuration	448
FIGURE 30-3	Example—PAP Authentication Scenario (Working From Home)	450
FIGURE 30-4	Example—CHAP Authentication Scenario (Calling a Private Network)	452
FIGURE 30-5	Example—PPPoE Tunnel	456
FIGURE 36-1	PAP Authentication Process	547
FIGURE 36-2	CHAP Authentication Sequence	550

Examples

EXAMPLE 16-1	Sample <code>/etc/auto_master</code> File	232
EXAMPLE 35-1	Output From a Properly Operating Dial-up Link	507
EXAMPLE 35-2	Output From a Properly Operating Leased-Line Link	508
EXAMPLE 35-3	Log File for a Link With a PPPoE Tunnel	519
EXAMPLE 35-4	PPPoE Diagnostic Messages	520
EXAMPLE 35-5	PPPoE snoop trace	521
EXAMPLE 36-1	To plumb an Interface to Support PPPoE	556
EXAMPLE 36-2	To List All Interfaces on a PPPoE Access Server	556
EXAMPLE 36-3	To Unplumb an Interface With a PPPoE Tunnel	557
EXAMPLE 36-4	Basic <code>/etc/ppp/pppoe</code> File	558
EXAMPLE 36-5	<code>/etc/ppp/pppoe</code> File for an Access Server	560
EXAMPLE 36-6	<code>/etc/ppp/options</code> File for an Access Server	561
EXAMPLE 36-7	<code>/etc/hosts</code> File for an Access Server	561
EXAMPLE 36-8	<code>/etc/ppp/pap-secrets</code> File for an Access Server	561
EXAMPLE 36-9	<code>/etc/ppp/chap-secrets</code> File for an Access Server	562
EXAMPLE 36-10	<code>/etc/ppp/peers/peer-name</code> to Define a Remote Access Server	563
EXAMPLE 40-1	Fields in <code>/etc/uucp/Systems</code>	588
EXAMPLE 40-2	Type Field and <code>/etc/uucp/Devices</code> File	589
EXAMPLE 40-3	Speed Field and <code>/etc/uucp/Devices</code> File	589
EXAMPLE 40-4	Phone Field Correspondence	590
EXAMPLE 40-5	Type Field and <code>/etc/uucp/Systems</code> File Equivalent	595
EXAMPLE 40-6	UUCP Class Field	595
EXAMPLE 40-7	Dialers Field for Direct Connect Modem	597
EXAMPLE 40-8	UUCP Dialers Field for Computers on Same Port Selector	597
EXAMPLE 40-9	UUCP Dialers Field for Modems Connected to Port Selector	597
EXAMPLE 40-10	<code>/etc/uucp/Dialers</code> File Entry	599

Composed September 5, 2002

EXAMPLE 40-11 Excerpts From `/etc/uucp/Dialers` 600

Preface

System Administration Guide: Resource Management and Network Services is part of a multivolume set that covers a significant part of the Solaris™ system administration information. 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, such as the Solaris Common Desktop Environment (CDE).

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 that run the Solaris 9 release. To use this book, you should have one to two 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

Related Books

This is a list of related documentation that is referred to in this book.

- *System Administration Guide: Advanced Administration*
- *System Administration Guide: Basic Administration*
- *System Administration Guide: IP Services*
- *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*
- *System Administration Guide: Naming and Directory Services (FNS and NIS+)*
- *System Administration Guide: Resource Management and Network Services*

- *System Administration Guide: Security Services*
- Anderson, Bart, Bryan Costales, and Harry Henderson. *UNIX Communications*. Howard W. Sams & Company, 1987.
- Costales, Bryan. *sendmail, Second Edition*. O'Reilly & Associates, Inc., 1997.
- Frey, Donnalyn and Rick Adams. *!%@:: A Directory of Electronic Mail Addressing and Networks*. O'Reilly & Associates, Inc., 1993.
- Krol, Ed. *The Whole Internet User's Guide and Catalog*. O' Reilly & Associates, Inc., 1993.
- O' Reilly, Tim and Grace Todino. *Managing UUCP and Usenet*. O' Reilly & Associates, Inc., 1992.

Related Information

For information on PPPoE licensing terms, refer to the incorporated material at the following locations:

`/var/sadm/pkg/SUNWpppd/install/copyright`

`/var/sadm/pkg/SUNWpppdu/install/copyright`

`/var/sadm/pkg/SUNWpppg/install/copyright`

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`.

Typographic Conventions

The following table describes the typographic changes 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% su</code> 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. You must be <i>root</i> to do this.

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>

CHAPTER **1**

Resource Management and Network Service (Overview)

Topics for the Solaris 9 Release

The following services or utilities are covered in this book:

Chapter 5	Resource management functionality enhances your ability to allocate, monitor, and control system resources.
Chapter 14	NFS is a protocol that provides the ability to access file systems from a remote host.
Chapter 18	SLP is a dynamic service discovery protocol.
Chapter 24	Mail services allow for a message to be sent to one or more people while routing the message over whatever networks are necessary.
Chapter 29	PPP is a protocol that provides point-to-point links between remote hosts.
Chapter 38	UUCP enables hosts to exchange files
Chapter 42	These commands are used to access files on remote systems. The commands include <code>ftp</code> , <code>rlogin</code> and <code>rcp</code> .
Chapter 2	NCA provides improved web server performance by caching web pages.
"Perl 5" on page 40	The Practical Extraction and Report Language (Perl) is a tool that can be used to generate scripts to assist with system administration tasks.

Perl 5

This Solaris release includes Practical Extraction and Report Language (Perl) 5.6.1, a powerful general-purpose programming language that is generally available as free software. Perl has emerged as the standard development tool for complex system administration task because of its excellent process, file, and text manipulation features.

Perl 5 includes a dynamically loadable module framework, which allows the addition of new capabilities for specific tasks. Many modules are freely available from the Comprehensive Perl Archive Network (CPAN) at <http://www.cpan.org>.

Accessing Perl Documentation

Several sources of information about Perl are included in this Solaris release. The same information is available by using these two mechanisms.

You can access the man pages by adding `/usr/perl5/man` to your `MANPATH` environment variable. This example displays the Perl overview.

```
% set MANPATH=($MANPATH /usr/perl5/man)
% man perl
```

You can access additional documentation by using the `perldoc` utility. This example displays the same overview information.

```
% /usr/perl5/bin/perldoc perl
```

The `perl` overview page lists of all the documentation that is included with the release.

Perl Compatibility Issues

In general, the 5.6.1 version of Perl is compatible with the previous version. Scripts do not have to be rebuilt or recompiled to function. However, any XSUB-based (.xs) modules require recompilation and reinstallation.

In the Solaris 9 release, you can access the older version of Perl as `/usr/perl5/5.00503/bin/perl`. The older version might not be supported in future releases and this version should only be used until the new modules are rebuilt.

Changes to the Solaris Version of Perl

The Solaris version of Perl was compiled to include system malloc, 64-bit integer and large file support. In addition, appropriate patches have been applied. For a full list of all configuration information, review the results from this command.

```
% /usr/perl5/bin/perlbug -dv
---
Flags:
    category=
    severity=
---
Site configuration information for perl v5.6.1:
.
.
```

You can generate a shorter list by using `perl -V`.

Composed September 5, 2002

Managing Web Cache Servers

This chapter provides an overview of the Solaris Network Cache and Accelerator (NCA). Also, procedures for using NCA and reference material about NCA are included.

- “Network Cache and Accelerator (Overview)” on page 43
- “Managing Web Cache Servers (Task Map)” on page 44
- “Administering the Caching of Web Pages (Tasks)” on page 46
- “Caching Web Pages (Reference)” on page 50

Network Cache and Accelerator (Overview)

The Solaris Network Cache and Accelerator (NCA) increases web server performance by maintaining an in-kernel cache of web pages that are accessed during HTTP requests. This in-kernel cache uses system memory to significantly increase performance for HTTP requests that are normally handled by web servers. Using system memory to hold web pages for HTTP requests increases web server performance by reducing the overhead between the kernel and the web server. NCA provides a sockets interface through which any web server can communicate with NCA with minimal modifications.

In situations where the requested page is retrieved from the in-kernel cache (cache hit), performance improves dramatically. In situations where the requested page is not in the cache (cache miss) and must be retrieved from the web server, performance is also significantly improved.

This product is intended to be run on a dedicated web server. If you run other large processes on a server that runs NCA, problems can result.

NCA provides logging support in that NCA logs all cache hits. This log is stored in binary format to increase performance. The `ncab2clf` command can be used to convert the log from binary format to common log format (CLF).

The Solaris 9 release includes the following enhancements:

- Sockets interface.
- Support for vectored sendfile, which provides support for AF_NCA. See the `sendfilev(3EXT)` man page for more information.
- New options for the `ncab2clf` command that support the ability to skip records before a selected date (`-s`) and to process a specified number of records (`-n`).
- `logd_path_name` in `ncalogd.conf` can specify either a raw device, a file, or a combination of the two.

New NCA Functionality

This release now provides support for a web server to open multiple AF-NCA sockets. With multiple sockets, you can have different web servers that run on one server.

In addition, a new configuration file that is called `/etc/nca/ncaport.conf` is available. The file can be used to manage the IP addresses and ports that NCA uses. Your web server might not provide native support of the AF_NCA socket. If your server lacks this support, use this file and the NCA socket utility library to convert an AF_INET socket to an AF_NCA socket.

Managing Web Cache Servers (Task Map)

The following table describes the procedures that are needed to use NCA.

Task	Description	For Instructions
Planning for NCA	A list of issues to be resolved before you enable the use of NCA.	"Planning for NCA" on page 45
Enabling NCA	Steps to enable in-kernel caching of web pages on a web server.	"How to Enable Caching of Web Pages" on page 46
Disabling NCA	Steps to disable in-kernel caching of web pages on a web server.	"How to Disable Caching of Web Pages" on page 48

Task	Description	For Instructions
Administering NCA logging	Steps to enable or disable the NCA logging process.	"How to Enable or Disable NCA Logging" on page 49
Loading the NCA socket library	Steps to use NCA if the AF_NCA socket is not supported.	"How to Load the Socket Utility Library for NCA" on page 49

Planning for NCA

The following sections cover the issues that need to be resolved before starting the NCA service.

System Requirements for NCA

To support NCA, the system must meet these requirements:

- 256 Mbytes RAM must be installed.
- The Solaris 9 release or one of the Solaris 8 upgrade releases must be installed.
- Apache support must be available. Apache support is available in the Solaris 9 and the Solaris 8 upgrade releases.

This product is intended to be run on a dedicated web server. The running of other large processes on a server that runs NCA can cause problems.

NCA Logging

The NCA service can be configured to log web activity. Generally, NCA logging should be enabled if the web server logging is enabled.

Interpositioning Library for Daemon Support of the Door Server

Many web servers use AF_INET sockets. By default, NCA uses AF_NCA sockets. To correct this situation, an interpositioning library is provided. The new library is loaded in front of the standard socket library, `libsocket.so`. The library call `bind()` is interposed by the new library, `ncad_addr.so`. Suppose that the status is enabled in `/etc/nca/ncakmod.conf`. The version of Apache that is included with the Solaris 9 release is already set up to call this library. If you are using IWS or Netscape servers, see "How to Load the Socket Utility Library for NCA" on page 49 to use the new library.

Multiple Instance Support

Systems that have NCA installed often need to run multiple instances of a web server. For instance, a single server might need to support a web server for outside access as well as a web administration server. To separate these servers, you would configure each server to use a separate port.

Administering the Caching of Web Pages (Tasks)

The following sections cover the procedures to enable or disable parts of the service.

▼ How to Enable Caching of Web Pages

1. Become superuser.

2. Register the interfaces.

Type the names of each of the physical interfaces in the `/etc/nca/nca.if` file. See the `nca.if(4)` man page for more information.

```
# cat /etc/nca/nca.if
hme0
hme1
```

Each interface must have an accompanying `hostname.interface-name` file and an entry in `/etc/hosts` file for the contents of `hostname.interface-name`. To start the NCA feature on all interfaces, place an asterisk, `*`, in the `nca.if` file.

3. Enable the `ncakmod` kernel module.

Change the status entry in `/etc/nca/ncakmod.conf` to enabled.

```
# cat /etc/nca/ncakmod.conf
#
# NCA Kernel Module Configuration File
#
status=enabled
httpd_door_path=/var/run/nca_httpd_1.door
nca_active=disabled
```

See the `ncakmod.conf(4)` man page for more information.

4. (Optional) Enable NCA logging.

Change the status entry in `/etc/nca/ncalogd.conf` to enabled.

```
# cat /etc/nca/ncalogd.conf
#
```

```
# NCA Logging Configuration File
#
status=enabled
logd_path_name="/var/nca/log"
logd_file_size=1000000
```

You can change the location of the log file by changing the path that is indicated by the `logd_path_name` entry. The log file can be a raw device or a file. See the following examples for samples of NCA log file paths. See the `ncalogd.conf(4)` man page for more information about the configuration file.

5. (Optional) Define ports for multiple instance support.

Add the port numbers in the `/etc/nca/ncaport.conf` file. This entry causes NCA to monitor port 80 on all configured IP addresses.

```
# cat /etc/nca/ncaport.conf
#
# NCA Kernel Module Port Configuration File
#
.
.
ncaport=*/80
```

6. For IA only: Increase the virtual memory size.

Use the `eeeprom` command to set the `kernelbase` of the system.

```
# eeeprom kernelbase=0x90000000
# eeeprom kernelbase
kernelbase=0x90000000
```

The second command verifies that the parameter has been set.

Note – By setting the `kernelbase`, you reduce the amount of virtual memory that user processes can use to less than 3 Gbytes. This restriction means that the system is not ABI compliant. When the system boots, the console displays a message that warns you about noncompliance. Most programs do not actually need the full 3-Gbyte virtual address space. If you have a program that needs more than 3 Gbytes, you need to run the program on a system that does not have NCA enabled.

7. Reboot the server.

Example—Using a Raw Device as the NCA Log File

The `logd_path_name` string in `ncalogd.conf` can define a raw device as the place to store the NCA log file. The advantage to using a raw device is that the service can run faster because the overhead in accessing a raw device is less.

The NCA service tests any raw device that is listed in the file to ensure that no file system is in place. This test ensures that no active file systems are accidentally written over.

To prevent this test from finding a file system, run the following command. This command destroys part of the file system on any disk partition that had been configured as a file system. In this example, `/dev/rdisk/c0t0d0s7` is the raw device that has an old file system in place.

```
# dd if=/dev/zero of=/dev/rdisk/c0t0d0s7 bs=1024 count=1
```

After running `dd`, you can then add the raw device to the `ncalogd.conf` file.

```
# cat /etc/nca/ncalogd.conf
#
# NCA Logging Configuration File
#
status=enabled
logd_path_name="/dev/rdisk/c0t0d0s7"
logd_file_size=1000000
```

Example—Using Multiple Files for NCA Logging

The `logd_path_name` string in `ncalogd.conf` can define multiple targets as the place to store the NCA log file. The second file is used when the first file is full. The following example shows how to select to write to the `/var/nca/log` file first and then use a raw partition.

```
# cat /etc/nca/ncalogd.conf
#
# NCA Logging Configuration File
#
status=enabled
logd_path_name="/var/nca/log /dev/rdisk/c0t0d0s7"
logd_file_size=1000000
```

▼ How to Disable Caching of Web Pages

1. Become superuser.

2. Disable the `ncakmod` kernel module.

Change the `status` entry in `/etc/nca/ncakmod.conf` to `disabled`.

```
# cat /etc/nca/ncakmod.conf
# NCA Kernel Module Configuration File
#
status=disabled
httpd_door_path=/var/run/nca_httpd_1.door
nca_active=disabled
```


See the `ncakmod.conf(4)` man page for more information.

3. Disable NCA logging.

Change the status entry in `/etc/nca/ncalogd.conf` to disabled.

```
# cat /etc/nca/ncalogd.conf
#
# NCA Logging Configuration File
#
status=disabled
logd_path_name="/var/nca/log"
logd_file_size=1000000
```

See the `ncalogd.conf(4)` man page for more information.

4. Reboot the server.

▼ How to Enable or Disable NCA Logging

NCA logging can be turned on or turned off, as needed, after NCA has been enabled. See “How to Enable Caching of Web Pages” on page 46 for more information.

1. Become superuser.

2. Change NCA logging.

To permanently disable logging, you need to change the status in `/etc/nca/ncalogd.conf` to disabled and reboot the system. See the `ncalogd.conf(4)` man page for more information.

a. Stop logging.

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

b. Start logging.

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

▼ How to Load the Socket Utility Library for NCA

Follow this process only if your web server does not provide native support of the AF_NCA socket.

In the startup script for the web server, add a line that causes the library to be preloaded. The line should resemble the following:

```
LD_PRELOAD=/usr/lib/ncad_addr.so /usr/bin/httpd
```

▼ How to Add a New Port to the NCA Service

1. **Become superuser.**

2. **Add a new port.**

Add a new port entry to `/etc/nca/ncaport.conf`. This example adds port 8888 on IP address 192.168.84.71. See `ncaport.conf(4)` for more information.

```
# cat /etc/nca/ncaport.conf
#
# NCA Kernel Module Port Configuration File
#
.
.
ncaport=*/80
ncaport=192.168.84.71/8888
```

3. **Start a new web instance.**

An address needs to be in the file that contains the NCA port configurations before a web server can use the address for NCA. If the web server is running, it must be restarted after the new address is defined.



Caching Web Pages (Reference)

The following sections cover the files and the components that are needed to use NCA. Also, specifics about how NCA interacts with the web server are included.

NCA Files

You need several files to support the NCA feature. Many of these files are ASCII, but some of the files are binary. The following table lists all of the files.

TABLE 2-1 NCA Files

File Name	Function
<code>/dev/nca</code>	The path name for the NCA device.
<code>/etc/hostname.*</code>	File that lists all physical interfaces that are configured on the server.

TABLE 2-1 NCA Files *(Continued)*

File Name	Function
<code>/etc/hosts</code>	File that lists all host names that are associated with the server. Entries in this file must match entries in <code>/etc/hostname.*</code> files for NCA to function.
<code>/etc/init.d/ncakmod</code>	Script that starts the NCA server. This script is run when a server is booted.
<code>/etc/init.d/ncalogd</code>	Script that starts NCA logging. This script is run when a server is booted.
<code>/etc/nca/nca.if</code>	File that lists the interfaces on which NCA is run. See the <code>nca.if(4)</code> man page for more information.
<code>/etc/nca/ncakmod.conf</code>	File that lists configuration parameters for NCA. See the <code>ncakmod.conf(4)</code> man page for more information.
<code>/etc/nca/ncalogd.conf</code>	File that lists configuration parameters for NCA logging. See the <code>ncalogd.conf(4)</code> man page for more information.
<code>/etc/nca/ncaport.conf</code>	File that lists the IP addresses and the ports for NCA. See the <code>ncaport.conf(4)</code> man page for more information.
<code>/usr/bin/ncab2clf</code>	Command that is used to convert data in the log file to the common log format. See the <code>ncab2clf(1)</code> man page for more information.
<code>/usr/lib/net/ncaconfd</code>	Command that is used to configure NCA to run on multiple interfaces during boot. See the <code>ncaconfd(1M)</code> man page for more information.
<code>/usr/lib/nca_addr.so</code>	Library that uses AF_NCA sockets instead of AF_INET sockets. This library must be used on web servers that use AF_INET sockets. See the <code>ncad_addr(4)</code> man page for more information.
<code>/var/nca/log</code>	File that holds the log file data. The file is in binary format, so do not edit it.
<code>/var/run/nca_httpd_1.door</code>	The door path name.

NCA Architecture

The NCA feature includes the following components.

- Kernel module, `ncakmod`
- Web server, `httpd`

The kernel module `ncakmod` maintains the cache of web pages in system memory. The module communicates with a web server, `httpd`, through a sockets interface. The family type is `PF_NCA`.

The kernel module also provides a logging facility that logs all HTTP cache hits. NCA logging writes HTTP data to the disk in binary format. NCA provides a conversion utility for converting binary log files to common log format (CLF).

The following figure shows the flow of data for the conventional path and the path that is used when NCA is enabled.




FIGURE 2-1 Data Flow With the NCA Service

NCA to Httpd Request Flow

The following list shows the request flow between the client and the web server.

1. An HTTP request is made from the client to the web server.
2. If the page is in cache, the in-kernel cache web page is returned.
3. If the page is not in cache, the request goes to the web server to retrieve or update the page.

Composed September 5, 2002

4. Depending on the HTTP protocol semantics that are used in the response, the page is cached or not. Then the page is returned to the client. If the Pragma: No-cache header is included in the HTTP request, the page is not cached.

Composed September 5, 2002

Time-Related Services

Keeping system clocks synchronized within a network is required for many databases and authentication services. The following topics are covered in this chapter.

- “Clock Synchronization (Overview)” on page 55
- “Managing Network Time Protocol (Tasks)” on page 56
- “Using Other Time-Related Commands (Tasks)” on page 57
- “Network Time Protocol (Reference)” on page 58

Clock Synchronization (Overview)

The Network Time Protocol (NTP) public domain software from the University of Delaware is included in the Solaris software from Solaris 2.6 release forward. The `xntpd` daemon sets and maintains the system time-of-day. The `xntpd` daemon is a complete implementation of the version 3 standard, as defined by RFC 1305.

The `xntpd` daemon reads the `/etc/inet/ntp.conf` file at system startup. See `xntpd(1M)` for information about configuration options.

Remember the following when using NTP in your network:

- The `xntpd` daemon uses minimal system resources.
- An NTP client synchronizes automatically with an NTP server when it boots. If the client becomes unsynchronized, the client resynchronizes again when the client contacts a time server.

Another way to synchronize clocks is to run `rdate` while using `cron`.

Managing Network Time Protocol (Tasks)

The following procedures show how to set up and use the NTP service.

▼ How to Set Up an NTP Server

1. **Become superuser.**

2. **Create the `ntp.conf` file.**

To ensure proper execution of the `xntpd` daemon, the `ntp.conf` file must first be created. The `ntp.server` file can be used as a template.

```
# cd /etc/inet
# cp ntp.server ntp.conf
```

3. **Start the `xntpd` daemon.**

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

▼ How to Set Up an NTP Client

1. **Become superuser.**

2. **Create the `ntp.conf` file.**

To activate the `xntpd` daemon, the `ntp.conf` file must first be created.

```
# cd /etc/inet
# cp ntp.client ntp.conf
```

3. **Start the `xntpd` daemon.**

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


Using Other Time-Related Commands (Tasks)

▼ How to Synchronize Date and Time From Another System

1. **Become superuser.**
2. **Reset the date and time to synchronize with another system, by using the `rdate` command.**

```
# rdate another-system
```

another-system Name of the another system

3. **Verify that you have reset your system's date correctly by using the `date` command.**
The output should show a date and time that matches that of the other system.

Example—Synchronizing Date and Time From Another System

The following example shows how to use `rdate` to synchronize the date and time of one system with another. In this example, the system `earth`, running several hours behind, is reset to match the date and time of the server `starbug`.

```
earth# date  
Tue Jun  5 11:08:27 MDT 2001  
earth# rdate starbug  
Tue Jun  5 14:06:37 2001  
earth# date  
Tue Jun  5 14:06:40 MDT 2001
```

Network Time Protocol (Reference)

The following files are needed for the NTP service to run.

TABLE 3-1 NTP Files

File Name	Function
<code>/etc/inet/ntp.conf</code>	Lists configuration options for NTP.
<code>/etc/inet/ntp.client</code>	Sample configuration file for NTP clients.
<code>/etc/inet/ntp.server</code>	Sample configuration file for NTP servers.
<code>/etc/inet/ntp.drift</code>	Sets the initial frequency offset on NTP servers.
<code>/etc/inet/ntp.keys</code>	Sample configuration file for NTP servers.
<code>/etc/init.d/xntpd</code>	NTP startup script run when a host is booted.
<code>/usr/lib/inet/xntpd</code>	NTP daemon. See <code>xntpd(1M)</code> for more information.
<code>/usr/sbin/ntpdate</code>	Utility to set the local date and time, based on NTP. See <code>ntpdate(1M)</code> for more information.
<code>/usr/sbin/ntpq</code>	NTP query program. See <code>ntpq(1M)</code> for more information.
<code>/usr/sbin/ntptrace</code>	Program to trace NTP hosts back to the master NTP server. See <code>ntptrace(1M)</code> for more information.
<code>/usr/sbin/xntpdcc</code>	NTP query program for the <code>xntpd</code> daemon. See <code>xntpdcc(1M)</code> for more information.
<code>/var/ntp/ntpstats</code>	Directory for holding NTP statistics.

Solaris 9 Resource Manager Topics

The section contains the following chapters on resource management in the Solaris operating environment.

Chapter 5	Provides an overview of resource management and discusses why you would want to use the functionality on your system
Chapter 6	Covers the project and task facilities and describes how they are used to label and separate workloads
Chapter 7	Describes the extended accounting functionality that is used to capture detailed resource consumption statistics for capacity planning or billing purposes
Chapter 8	Discusses resource controls, which are used to place bounds on resource usage by applications that run on your system
Chapter 9	Describes the fair share scheduler, which uses shares to specify the amounts of CPU time that is allocated to processes that run on your system
Chapter 10	Describes resource pools, which are used to partition system resources and guarantee that a known amount of resources is always available to a specified workload that runs on your system
Chapter 11	Describes a hypothetical server consolidation project

Composed September 5, 2002

Chapter 12	Describes the resource management functionality available in the Solaris Management Console tool

Introduction to Solaris 9 Resource Manager

Resource management functionality enables you to control how applications use available system resources. You can do the following:

- Allocate computing resources, such as processor time
- Monitor how the allocations are being used, then adjust the allocations as necessary
- Generate extended accounting information for analysis, billing, and capacity planning

Overview

Modern computing environments have to provide a flexible response to the varying workloads that are generated by different applications on a system. If resource management features are not used, the Solaris operating environment responds to workload demands by adapting to new application requests dynamically. This default response generally means that all activity on the system is given equal access to resources. Solaris resource management features enable you to treat workloads individually. You can do the following:

- Restrict access to a specific resource
- Offer resources to workloads on a preferential basis
- Isolate workloads from one another

The ability to minimize cross-workload performance compromises, along with the facilities that monitor resource usage and utilization, is referred to as *resource management*. Resource management is implemented through a collection of algorithms. The algorithms handle the series of capability requests that an application presents in the course of its execution.

Resource management facilities permit you to modify the default behavior of the operating system with respect to different workloads. *Behavior* primarily refers to the set of decisions that are made by operating system algorithms when an application presents one or more resource requests to the system. You can use resource management facilities to do the following:

- Deny resources or prefer one application over another for a larger set of allocations than otherwise permitted
- Treat certain allocations collectively instead of through isolated mechanisms

The implementation of a system configuration that uses the resource management facilities can serve several purposes. You can do the following:

- Prevent an application from consuming resources indiscriminately
- Change an application's priority based on external events
- Balance resource guarantees to a set of applications against the goal of maximizing system utilization

When planning a resource-managed configuration, key requirements include the following:

- Identifying the competing workloads on the system
- Distinguishing those workloads that are not in conflict from those workloads with performance requirements that compromise the primary workloads

After you identify cooperating and conflicting workloads, you can create a resource configuration that presents the least compromise to the service goals of the business, within the limitations of the system's capabilities.

Effective resource management is enabled in the Solaris environment by offering control mechanisms, notification mechanisms, and monitoring mechanisms. Control mechanisms are described in "Resource Management Control Mechanisms" on page 63. Notification mechanisms and monitoring mechanisms are described in Chapter 8. Many of these capabilities are provided through enhancements to existing mechanisms such as the `proc(4)` file system, processor sets, and scheduling classes. Other capabilities are specific to resource management. These capabilities are described in subsequent chapters.

Resource Classifications

A resource is any aspect of the computing system that can be manipulated with the intent to change application behavior. Thus, a resource is a capability that an application implicitly or explicitly requests. If the capability is denied or constrained, the execution of a robustly written application proceeds more slowly.

Classification of resources, as opposed to identification of resources, can be made along a number of axes. The axes could be implicitly requested as opposed to explicitly requested, time-based (such as CPU time) compared to time-independent (such as CPU shares assigned), and so forth.

Generally, scheduler-based resource management is applied to resources that the application can implicitly request. For example, to continue execution, an application implicitly requests additional CPU time. To write data to a network socket, an application implicitly requests bandwidth. Constraints can be placed on the aggregate total use of an implicitly requested resource.

Additional interfaces can be presented so that bandwidth or CPU service levels can be explicitly negotiated. Resources that are explicitly requested, such as a request for an additional thread, can be managed by constraint.

Resource Management Control Mechanisms

The three types of control mechanisms that are available in the Solaris operating environment are constraints, scheduling, and partitioning.

Constraints

Constraints allow the administrator or application developer to set bounds on the consumption of specific resources for a workload. With known bounds, modeling resource consumption scenarios becomes a simpler process. Bounds can also be used to control ill-behaved applications that would otherwise compromise system performance or availability through unregulated resource requests.

Constraints do present complications for the application. The relationship between the application and the system can be modified to the point that the application is no longer able to function. One approach that can mitigate this risk is to gradually narrow the constraints on applications with unknown resource behavior. The resource controls feature discussed in Chapter 8 provides a constraint mechanism. Newer applications can be written to be aware of their resource constraints, but not all application writers will choose to do this.

Scheduling

Scheduling refers to making a sequence of allocation decisions at specific intervals. The decision that is made is based on a predictable algorithm. An application that does not need its current allocation leaves the resource available for another application's use. Scheduling-based resource management enables full utilization of an undercommitted configuration, while providing controlled allocations in a critically committed or overcommitted scenario. The underlying algorithm defines how the

term “controlled” is interpreted. In some instances, the scheduling algorithm might guarantee that all applications have some access to the resource. The fair share scheduler (FSS) described in Chapter 9 manages application access to CPU resources in a controlled way.

Partitioning

Partitioning is used to bind a workload to a subset of the system’s available resources. This binding guarantees that a known amount of resources is always available to the workload. The resource pools functionality described in Chapter 10 enables you to limit workloads to specific subsets of the machine. Configurations that use partitioning can avoid system-wide overcommitment. However, in avoiding this overcommitment, the possibility of achieving high utilizations is less likely. A reserved group of resources (such as processors) is not available for use by another workload when the workload bound to them is idle.

Resource Management Configuration

Portions of the resource management configuration can be placed in a network name service. This feature allows the administrator to apply resource management constraints across a collection of machines, rather than on an exclusively per-machine basis. Related work can share a common identifier, and the aggregate usage of that work can be tabulated from accounting data.

Resource management configuration and workload-oriented identifiers are described more fully in Chapter 6. The extended accounting facility that links these identifiers with application resource usage is described in Chapter 7.

When to Use Resource Management

Use resource management to ensure that your applications have the response times they require.

Resource management can also increase resource utilization. By categorizing and prioritizing usage, you can effectively use reserve capacity during off-peak periods, often eliminating the need for additional processing power. You can also ensure that resources are not wasted because of load variability.

Server Consolidation

Resource management is ideal for environments that consolidate a number of applications on a single server.

The cost and complexity of managing numerous machines encourages the consolidation of several applications on larger, more scalable servers. Instead of running each workload on a separate system, with full access to that system's resources, you can use resource management software to segregate workloads within the system. Resource management enables you to lower overall total cost of ownership by running and controlling several dissimilar applications on a single Solaris system.

If you are providing Internet and application services, you can use resource management to do the following:

- Host multiple web servers on a single machine. You can control the resource consumption for each web site and protect each site from the potential excesses of other sites.
- Prevent a faulty common gateway interface (CGI) script from exhausting CPU resources.
- Stop an incorrectly behaving application from leaking all available virtual memory.
- Ensure that one customer's applications are not affected by another customer's applications that run at the same site.
- Provide differentiated levels or classes of service on the same machine.
- Obtain accounting information for billing purposes.

Supporting a Large or Varied User Population

Use resource management features in any system that has a large, diverse user base, such as an educational institution. If you have a mix of workloads, the software can be configured to give priority to specific projects.

For example, in large brokerage firms, traders intermittently require fast access to execute a query or to perform a calculation. Other system users, however, have more consistent workloads. If you allocate a proportionately larger amount of processing power to the traders' projects, the traders have the responsiveness they need.

Resource management is also ideal for supporting thin-client systems. These platforms provide stateless consoles with frame buffers and input devices, such as smart cards. The actual computation is done on a shared server, resulting in a timesharing type of environment. Use resource management features to isolate the users on the server. Then, a user who generates excess load does not monopolize hardware resources and significantly impact others who use the system.

Setting Up Resource Management (Task Map)

The following task map gives a basic overview of the steps that are involved in setting up resource management on your system.

Task	Description	For Instructions
Identify the workloads on your system.	Review project entries in either the <code>/etc/project</code> database file or in the NIS map or LDAP directory service.	"project Database" on page 70
Prioritize the workloads on your system.	Determine which applications are critical. These workloads might require preferential access to resources.	Refer to your business service goals.
Monitor real-time activity on your system.	Use performance tools to view the current resource consumption of workloads that are running on your system. You can then evaluate whether you must restrict access to a particular resource or isolate workloads from one another.	"Monitoring by System" on page 136, <code>cpustat(1M)</code> , <code>iostat(1M)</code> , <code>mpstat(1M)</code> , <code>prstat(1M)</code> , <code>sar(1)</code> , and <code>vmstat(1M)</code>
Make temporary modifications to the workloads that are running on your system.	To determine which values can be altered, refer to the resource controls that are available in the Solaris environment. You can update the values from the command line while the task or process is running.	"Available Resource Controls" on page 90, "Actions on Resource Control Values" on page 92, "Temporarily Updating Resource Control Values on a Running System" on page 95, <code>rctladm(1M)</code> , and <code>prctl(1)</code>

Composed September 5, 2002

Task	Description	For Instructions
Set resource control attributes for every project entry in the project database or name service project table.	Each project entry in the <code>/etc/project</code> database or the name service project table can contain one or more resource controls. The resource controls constrain tasks and processes attached to that project. For each threshold value that is placed on a resource control, you can associate one or more actions to be taken when that value is reached. You can set resource controls by using the command-line interface or the Solaris Management Console. If you are setting configuration parameters across a large number of systems, use the console for this task.	"project Database" on page 70, "Local project File Format" on page 71, "Available Resource Controls" on page 90, "Actions on Resource Control Values" on page 92, and Chapter 9
Create resource pools configurations.	Resource pools provide a way to partition system resources, such as processors, and maintain those partitions across reboots. You can add a <code>project.pool</code> attribute to each entry in the <code>/etc/project</code> database.	"Creating Pools Configurations" on page 119
Make the fair share scheduler (FSS) your default system scheduler.	Ensure that all user processes in either a single CPU system or a processor set belong to the same scheduling class.	"FSS Configuration Examples" on page 111 and <code>dispadm(1M)</code>
Activate the extended accounting facility to monitor and record resource consumption on a task or process basis.	Use extended accounting data to assess current resource controls and plan capacity requirements for future workloads. Aggregate usage on a system-wide basis can be tracked. To obtain complete usage statistics for related workloads that span more than one system, the project name can be shared across several machines.	"How to Activate Extended Accounting for Processes, Tasks, and Flows" on page 85 and <code>acctadm(1M)</code>

Task	Description	For Instructions
(Optional) If you determine that additional adjustments to your configuration are required, you can continue to alter the values from the command line while the task or process is running.	Modifications to existing tasks can be applied on a temporary basis without restarting the project. Tune the values until you are satisfied with the performance. Then update the current values in the <code>/etc/project</code> database or the name service project table.	"Temporarily Updating Resource Control Values on a Running System" on page 95, <code>rctladm(1M)</code> , and <code>prctl(1)</code>
(Optional) Capture extended accounting data.	Write extended accounting records for active processes and tasks. The files that are produced can be used for planning, chargeback, and billing purposes.	<code>wracct(1M)</code>

Projects and Tasks

This chapter discusses the *project* and *task* facilities of Solaris resource management. Projects and tasks are used to label workloads and separate them from one another. The project provides a network-wide administrative identifier for related work. The task collects a group of processes into a manageable entity that represents a workload component.

Overview

To optimize workload response, you must first be able to identify the workloads that are running on the system you are analyzing. This information can be difficult to obtain by using either a purely process-oriented or a user-oriented method alone. In the Solaris environment, you have two additional facilities that can be used to separate and identify workloads: the project and the task.

Based on their project or task membership, running processes can be manipulated with standard Solaris commands. The extended accounting facility can report on both process usage and task usage, and tag each record with the governing project identifier. This process enables offline workload analysis to be correlated with online monitoring. The project identifier can be shared across multiple machines through the project name service database. Thus, the resource consumption of related workloads that run on (or span) multiple machines can ultimately be analyzed across all of the machines.

Projects

The project identifier is an administrative identifier that is used to identify related work. The project identifier can be thought of as a workload tag equivalent to the user and group identifiers. A user or group can belong to one or more projects. These projects can be used to represent the workloads in which the user or group of users is allowed to participate. This membership can then be the basis of chargeback that is based on, for example, usage or initial resource allocations. Although a user must have a default project assigned, the processes that the user launches can be associated with any of the projects of which that user is a member.

Determining a User's Default Project

To log in to the system, a user must be assigned a default project.

Because each process on the system possesses project membership, an algorithm to assign a default project to the login or other initial process is necessary. The algorithm to determine a default project consists of four steps. If no default project is found, the user's login, or request to start a process, is denied.

The system sequentially follows these steps to determine a user's default project:

1. If the user has an entry with a `project` attribute defined in the `/etc/user_attr` extended user attributes database, then the value of the `project` attribute is the default project (see `user_attr(4)`).
2. If a project with the name `user.user-id` is present in the `project(4)` database, then that project is the default project.
3. If a project with the name `group.group-name` is present in the `project` database, where `group-name` is the name of the default group for the user (as specified in `passwd(4)`), then that project is the default project.
4. If the special project `default` is present in the `project` database, then that project is the default project.

This logic is provided by the `getdefaultproj()` library function (see `getproject(3PROJECT)`).

project Database

You can store project data in a local file, in a Network Information Service (NIS) project map, or in a Lightweight Directory Access Protocol (LDAP) directory service. The `/etc/project` database or name service is used at login and by all requests for account management by the pluggable authentication module (PAM) to bind a user to a default project.

Note – Updates to entries in the project database, whether to the `/etc/project` file or to a representation of the database in a network name service, are not applied to currently active projects. The updates are applied to new tasks that join the project when `login(1)` or `newtask(1)` is used.

PAM Subsystem

Operations that change or set identity include logging in to the system, invoking an `rcp` or `rsh` command, using `ftp`, or using `su`. When an operation involves changing or setting identity, a set of configurable modules is used to provide authentication, account management, credentials management, and session management.

The account management PAM module for projects is documented in the `pam_projects(5)` man page. The PAM system is documented in the man pages `pam(3PAM)`, `pam.conf(4)`, and `pam_unix(5)`.

Name Service Configuration

Resource management supports the name service `project` database. The location where the `project` database is stored is defined in `/etc/nsswitch.conf`. By default, `files` is listed first, but the sources can be listed in any order.

```
project: files [nis] [ldap]
```

If more than one source for project information is listed, the `nsswitch.conf` file directs the routine to start searching for the information in the first source listed. The routine then searches subsequent databases.

For more information on `/etc/nsswitch.conf`, see “The Name Service Switch (Overview)” in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)* and `nsswitch.conf(4)`.

Local project File Format

If you select `files` as your project database in `nsswitch.conf`, the login process searches the `/etc/project` file for project information (see `projects(1)` and `project(4)`). The `project` file contains a one-line entry for each project recognized by the system, of the following form:

```
projname:projid:comment:user-list:group-list:attributes
```

The fields are defined as follows.

<i>projname</i>	The name of the project. The name must be a string that consists of alphanumeric characters, the underline (_) character, and the hyphen (-). The name must begin with an alphabetic character. <i>projname</i> cannot contain periods (.), colons (:), or newline characters.
<i>projid</i>	The project's unique numerical ID (PROJID) within the system. The maximum value of the <i>projid</i> field is UID_MAX (2147483647).
<i>comment</i>	The project's description.
<i>user-list</i>	A comma-separated list of users who are allowed in the project. Wildcards can be used in this field. The asterisk (*) allows all users to join the project. The exclamation point followed by the asterisk (!*) excludes all users from the project. The exclamation mark (!) followed by a user name excludes the specified user from the project.
<i>group-list</i>	A comma-separated list of groups of users who are allowed in the project. Wildcards can be used in this field. The asterisk (*) allows all groups to join the project. The exclamation point followed by the asterisk (!*) excludes all groups from the project. The exclamation mark (!) followed by a group name excludes the specified group from the project.
<i>attributes</i>	A semicolon-separated list of name-value pairs (see Chapter 8). <i>name</i> is an arbitrary string that specifies the object-related attribute, and <i>value</i> is the optional value for that attribute. name [=value] In the name-value pair, names are restricted to letters, digits, underscores, and the period. The period is conventionally used as a separator between the categories and subcategories of the <code>rctl</code> . The first character of an attribute name must be a letter. The name is case sensitive. Values can be structured by using commas and parentheses to establish precedence. The semicolon is used to separate name-value pairs. The semicolon cannot be used in a value definition. The colon is used to separate project fields. The colon cannot be used in a value definition.

Note – Routines that read this file halt when they encounter a malformed entry. Any project assignments that are specified after the incorrect entry are not made.

This example shows the default `/etc/project` file:

```
system:0:System:::
user.root:1:Super-User:::
noproject:2:No Project:::
default:3:::
```



```
group.staff:10:::
```

This example shows the default `/etc/project` file with project entries added at the end:

```
system:0:System:::  
user.root:1:Super-User:::  
noproject:2:No Project:::  
default:3:::  
group.staff:10:::  
user.ml:2424:Lyle Personal:::  
booksite:4113:Book Auction Project:ml,mp,jtd,kjh::
```

To add resource controls to the `/etc/project` file, see “Using Resource Controls” on page 96.

Name Service Configuration for NIS

If you are using NIS, you can specify in the `/etc/nsswitch.conf` file to search the NIS maps for projects:

```
project: nis files
```

The NIS map, either `project.byname` or `project.bynumber`, has the same form as the `/etc/project` file:

```
projname:projid:comment:user-list:group-list:attributes
```

For more information, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Directory Service Configuration for LDAP

If you are using LDAP, you can specify in the `/etc/nsswitch.conf` file to search the LDAP entries for projects.

```
project: ldap files
```

For more information, including the schema for project entries in an LDAP database, see “General Reference (Reference)” in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Tasks

With each successful login into a project, a new *task* that contains the login process is created. The task is a process collective that represents a set of work over time. A task can also be viewed as a *workload component*.

Each process is a member of one task, and each task is associated with one project.




FIGURE 6-1 Project and Task Tree

All operations on sessions, such as signal delivery, are also supported on tasks. You can also bind tasks to processor sets and set their scheduling priorities and classes, which modifies all current and subsequent processes in the task.

Tasks are created at login (see `login(1)`), by `cron(1M)`, by `newtask(1)`, and by `setproject(3PROJECT)`.

The extended accounting facility can provide accounting data for processes that is aggregated at the task level.

Commands Used to Administer Projects and Tasks

Command	Description
<code>projects(1)</code>	Prints the project membership of a user.

Command	Description
<code>newtask(1)</code>	Executes the user's default shell or specified command, placing the execution command in a new task that is owned by the specified project. <code>newtask</code> can also be used to modify the task and the project binding for a running process.
<code>projadd(1M)</code>	Adds a new project entry to the <code>/etc/project</code> file. <code>projadd</code> creates a project entry only on the local system. <code>projadd</code> cannot change information that is supplied by the network name service.
<code>projmod(1M)</code>	Modifies a project's information on the local system. <code>projmod</code> cannot change information that is supplied by the network name service. However, the command does verify the uniqueness of the project name and project ID against the external name service.
<code>projdel(1M)</code>	Deletes a project from the local system. <code>projdel</code> cannot change information that is supplied by the network name service.

Command Options Used With Projects and Tasks

`ps`

Use `ps -o` to display task and project IDs. For example, to view the project ID, type the following:

```
# ps -o user,pid,uid,projid
USER PID  UID  PROJID
jtd  89430 124   4113
```

`id`

Use `id -p` to print the current project ID in addition to the user and group IDs. If the *user* operand is provided, the project associated with that user's normal login is printed:

```
# id -p
uid=124(jtd) gid=10(staff) projid=4113(booksite)
```

`pgrep` and `pkill`

To match only processes with a project ID in a specific list, type the following:

```
# pgrep -J projidlist
# pkill -J projidlist
```

To match only processes with a task ID in a specific list, type the following:

```
# pgrep -T taskidlist
# pkill -T taskidlist
```

prstat

To display various statistics for processes and projects that are currently running on your system, type the following:

```
% prstat -J
      PID USERNAME  SIZE  RSS STATE PRI NICE      TIME  CPU PROCESS/NLWP
21634 jtd          5512K 4848K cpu0   44    0  0:00.00 0.3% prstat/1
   324 root           29M   75M sleep   59    0  0:08.27 0.2% Xsun/1
15497 jtd           48M   41M sleep   49    0  0:08.26 0.1% adeptedit/1
   328 root        2856K 2600K sleep   58    0  0:00.00 0.0% mibiisa/11
  1979 jtd        1568K 1352K sleep   49    0  0:00.00 0.0% csh/1
  1977 jtd        7256K 5512K sleep   49    0  0:00.00 0.0% dtterm/1
   192 root        3680K 2856K sleep   58    0  0:00.36 0.0% automountd/5
  1845 jtd           24M   22M sleep   49    0  0:00.29 0.0% dtmail/11
  1009 jtd        9864K 8384K sleep   49    0  0:00.59 0.0% dtwm/8
   114 root        1640K  704K sleep   58    0  0:01.16 0.0% in.routed/1
   180 daemon      2704K 1944K sleep   58    0  0:00.00 0.0% statd/4
   145 root        2120K 1520K sleep   58    0  0:00.00 0.0% ypbind/1
   181 root        1864K 1336K sleep   51    0  0:00.00 0.0% lockd/1
   173 root        2584K 2136K sleep   58    0  0:00.00 0.0% inetd/1
   135 root        2960K 1424K sleep    0    0  0:00.00 0.0% keyserv/4
PROJID  NPROC  SIZE  RSS MEMORY      TIME  CPU PROJECT
   10     52  400M  271M   68%  0:11.45 0.4% booksite
    0     35  113M  129M   32%  0:10.46 0.2% system
```

Total: 87 processes, 205 lwps, load averages: 0.05, 0.02, 0.02

To display various statistics for processes and tasks that are currently running on your system, type the following:

```
% prstat -T
      PID USERNAME  SIZE  RSS STATE PRI NICE      TIME  CPU PROCESS/NLWP
23023 root           26M   20M sleep   59    0  0:03:18 0.6% Xsun/1
23476 jtd           51M   45M sleep   49    0  0:04:31 0.5% adeptedit/1
23432 jtd        6928K 5064K sleep   59    0  0:00:00 0.1% dtterm/1
28959 jtd           26M   18M sleep   49    0  0:00:18 0.0% .netscape.bin/1
23116 jtd        9232K 8104K sleep   59    0  0:00:27 0.0% dtwm/5
29010 jtd        5144K 4664K cpu0   59    0  0:00:00 0.0% prstat/1
   200 root        3096K 1024K sleep   59    0  0:00:00 0.0% lpsched/1
   161 root        2120K 1600K sleep   59    0  0:00:00 0.0% lockd/2
   170 root        5888K 4248K sleep   59    0  0:03:10 0.0% automountd/3
   132 root        2120K 1408K sleep   59    0  0:00:00 0.0% ypbind/1
   162 daemon      2504K 1936K sleep   59    0  0:00:00 0.0% statd/2
   146 root        2560K 2008K sleep   59    0  0:00:00 0.0% inetd/1
   122 root        2336K 1264K sleep   59    0  0:00:00 0.0% keyserv/2
   119 root        2336K 1496K sleep   59    0  0:00:02 0.0% rpcbind/1
   104 root        1664K  672K sleep   59    0  0:00:03 0.0% in.rdisc/1
```

Composed September 5, 2002

TASKID	NPROC	SIZE	RSS	MEMORY	TIME	CPU	PROJECT
222	30	229M	161M	44%	0:05:54	0.6%	group.staff
223	1	26M	20M	5.3%	0:03:18	0.6%	group.staff
12	1	61M	33M	8.9%	0:00:31	0.0%	group.staff
1	33	85M	53M	14%	0:03:33	0.0%	system

Total: 65 processes, 154 lwps, load averages: 0.04, 0.05, 0.06

Note – The `-J` and `-T` options cannot be used together.

Using cron and su With Projects and Tasks

cron

The `cron` command issues a `settaskid` to ensure that each `cron`, `at`, and `batch` job executes in a separate task, with the appropriate default project for the submitting user. Also, the `at` and `batch` commands capture the current project ID and ensure that the project ID is restored when running an `at` job.

su

To switch the user's default project, and thus create a new task (as part of simulating a login) type the following:

```
# su - user
```

To retain the project ID of the invoker, issue `su` without the `-` flag.

```
# su user
```

Project Administration Examples

▼ How to Define a Project and View the Current Project

This example shows how to use the `projadd` and `projmod` commands.

1. **Become superuser.**
2. **View the default `/etc/project` file on your system.**

```
# cat /etc/project
system:0:::
user.root:1:::
noproject:2:::
default:3:::
group.staff:10:::
```

3. Add a project called *booksite* and assign it to a user named *mark* with project ID number 4113.

```
# projadd -U mark -p 4113 booksite
```

4. View the `/etc/project` file again to see the project addition.

```
# cat /etc/project
system:0:::
user.root:1:::
noproject:2:::
default:3:::
group.staff:10:::
booksite:4113::mark::
```

5. Add a comment that describes the project in the comment field.

```
# projmod -c 'Book Auction Project' booksite
```

6. View the changes in the `/etc/project` file.

```
# cat /etc/project
system:0:::
user.root:1:::
noproject:2:::
default:3:::
group.staff:10:::
booksite:4113:Book Auction Project:mark::
```

▼ How to Delete a Project From the `/etc/project` File

This example shows how to use the `projdel` command to delete a project.

1. Become superuser.
2. Remove the project *booksite* by using the `projdel` command.

```
# projdel booksite
```

3. Display the `/etc/project` file.

```
# cat /etc/project
system:0:::
user.root:1:::
```

```
noproject:2:::  
default:3:::  
group.staff:10:::
```

4. Log in as user *mark* and type `projects` to view the projects assigned.

```
# su - mark  
# projects  
default
```

▼ How to Obtain User and Project Membership Information

Use the `id` command with the `-p` flag to view the current project membership of the invoking process.

```
$ id -p  
uid=100(mark) gid=1(other) projid=3(default)
```

▼ How to Create a New Task

1. Become superuser.
2. Create a new task in the *booksite* project by using the `newtask` command with the `-v` (verbose) option to obtain the system task ID.

```
# newtask -v -p booksite  
16
```

The execution of `newtask` creates a new task in the specified project, and places the user's default shell in this task.

3. View the current project membership of the invoking process.

```
# id -p  
uid=100(mark) gid=1(other) projid=4113(booksite)
```

The process is now a member of the new project.

▼ How to Move a Running Process Into a New Task

This example shows how to associate a running process with a different task and project. To perform this task, you must either be superuser, or be the owner of the process and be a member of the new project.

1. Become superuser.
2. Obtain the process ID of the *book_catalog* process.

Composed September 5, 2002

```
# pgrep book_catalog
8100
```

3. Associate process 8100 with a new task ID in the *booksite* project.

```
# newtask -v -p booksite -c 8100
17
```

The `-c` option specifies that `newtask` operate on the existing named process.

4. Confirm the task to process ID mapping.

```
# pgrep -T 17
8100
```


Extended Accounting

By using the project and task facilities that are described in Chapter 6 to label and separate workloads, you can monitor resource consumption by each workload. You can use the *extended accounting* subsystem to capture a detailed set of resource consumption statistics on both running processes and tasks. The extended accounting subsystem labels the usage records with the project for which the work was done. You can also use extended accounting, in conjunction with the Internet Protocol Quality of Service (IPQoS) flow accounting module described in “Using Flow Accounting and Statistics Gathering (Tasks)” in the *IPQoS Administration Guide*, to capture network flow information on a system.

To begin using extended accounting, see “How to Activate Extended Accounting for Processes, Tasks, and Flows” on page 85.

Overview

Before you can apply resource management mechanisms, you must first be able to characterize the resource consumption demands that various workloads place on a system. The extended accounting facility in the Solaris operating environment provides a flexible way to record system and network resource consumption on a task or process basis, or on the basis of selectors provided by IPQoS (see `ipqos(7IPP)`). Unlike online monitoring tools, which measure system usage in real time, extended accounting enables you to examine historical usage. You can then make assessments of capacity requirements for future workloads.

With extended accounting data available, you can develop or purchase software for resource chargeback, workload monitoring, or capacity planning.

How Extended Accounting Works

The extended accounting facility in the Solaris environment uses a versioned, extensible file format to contain accounting data. Files that use this data format can be accessed or be created by using the API provided in the included library, `libexacct(3LIB)`. These files can then be analyzed on any platform with extended accounting enabled, and their data can be used for capacity planning and chargeback.

If extended accounting is active, statistics are gathered that can be examined by the `libexacct` API. `libexacct` allows examination of the `exacct` files either forward or backward. The API supports third-party files that are generated by `libexacct` as well as those files that are created by the kernel.

With extended accounting enabled, the task tracks the aggregate resource usage of its member processes. A task accounting record is written at task completion. Interim records can also be written. For more information on tasks, see Chapter 6.




FIGURE 7-1 Task Tracking With Extended Accounting Activated

Extensible Format

The extended accounting format is substantially more extensible than the SunOS™ legacy system accounting software format (see “What is System Accounting?” in *System Administration Guide: Advanced Administration*). Extended accounting permits accounting metrics to be added and removed from the system between releases, and even during system operation.

Note – Both extended accounting and legacy system accounting software can be active on your system at the same time.

exacct Records and Format

Routines that allow `exacct` records to be created serve two purposes.

- To enable third-party `exacct` files to be created
- To enable the creation of tagging records to be embedded in the kernel accounting file by using the `putacct` system call (see `getacct(2)`)

The format permits different forms of accounting records to be captured without requiring that every change be an explicit version change. Well-written applications that consume accounting data must ignore records they do not understand.

The `libexacct` library converts and produces files in the `exacct` format. This library is the *only* supported interface to `exacct` format files.

Note – The `getacct`, `putacct`, and `wracct` system calls do not apply to flows. The kernel creates flow records and writes them to the file when IPQoS flow accounting is configured.

Extended Accounting Configuration

The `/etc/acctadm.conf` file contains the current extended accounting configuration. The file is edited through the `acctadm` interface, not by the user.

The directory `/var/adm/exacct` is the standard location for placing extended accounting data. You can use the `acctadm(1M)` command to specify a different location for the process and task accounting-data files.

Commands Used With Extended Accounting

Command	Description
<code>acctadm(1M)</code>	Modifies various attributes of the extended accounting facility, stops and starts extended accounting, and is used to select accounting attributes to track for processes, tasks, and flows.
<code>wracct(1M)</code>	Writes extended accounting records for active processes and active tasks.
<code>lastcomm(1)</code>	Displays previously invoked commands. <code>lastcomm</code> can consume either standard accounting-process data or extended-accounting process data.

For information on commands that are associated with tasks and projects, see “Commands Used to Administer Projects and Tasks” on page 74. For information on IPQoS flow accounting, see `ipqosconf(1M)`.

Using Extended Accounting Functionality

▼ How to Activate Extended Accounting for Processes, Tasks, and Flows

To activate the extended accounting facility for tasks, processes, and flows, use the `acctadm(1M)` command. The optional final parameter to `acctadm` indicates whether the command should act on the process, system task, or flow accounting components of the extended accounting facility.

1. Become superuser.

2. Activate extended accounting for processes.

```
# acctadm -e extended -f /var/adm/exacct/proc process
```

3. Activate extended accounting for tasks.

```
# acctadm -e extended,mstate -f /var/adm/exacct/task task
```

4. Activate extended accounting for flows.

```
# acctadm -e extended -f /var/adm/exacct/flow flow
```

▼ How to Activate Extended Accounting With a Startup Script

Activate extended accounting on an ongoing basis by linking the `/etc/init.d/acctadm` script into `/etc/rc2.d`.

```
# ln -s /etc/init.d/acctadm /etc/rc2.d/Snacctadm
# ln -s /etc/init.d/acctadm /etc/rc2.d/Knacctadm
```

The *n* variable is replaced by a number.

See “Extended Accounting Configuration” on page 84 for information on accounting configuration.

▼ How to Display Extended Accounting Status

Type `acctadm` without arguments to display the current status of the extended accounting facility.

```
# acctadm
    Task accounting: active
    Task accounting file: /var/adm/exacct/task
    Tracked task resources: extended
    Untracked task resources: none
    Process accounting: active
    Process accounting file: /var/adm/exacct/proc
    Tracked process resources: extended
    Untracked process resources: host,mstate
    Flow accounting: active
    Flow accounting file: /var/adm/exacct/flow
    Tracked flow resources: extended
    Untracked flow resources: none
```

In the previous example, system task accounting is active in extended mode and mstate mode. Process and flow accounting are active in extended mode.

Note – In the context of extended accounting, microstate (`mstate`) refers to the extended data, associated with microstate process transitions, that is available in the process usage file (see `proc(4)`). This data provides much more detail about the activities of the process than basic or extended records.

▼ How to View Available Accounting Resources

Available resources can vary from system to system, and from platform to platform. Use the `-r` option to view the available accounting resources on the system.

```
# acctadm -r
process:
extended pid,uid,gid,cpu,time,command,TTY,projid,taskid,ancpid,wait-status,flag
basic    pid,uid,gid,cpu,time,command,TTY,flag
task:
extended taskid,projid,cpu,time,host,mstate,anctaskid
basic    taskid,projid,cpu,timeprocess:
extended pid,uid,gid,cpu,time,command,TTY,projid,taskid,ancpid,wait-status,flag
basic    pid,uid,gid,cpu,time,command,TTY,flag
task:
extended taskid,projid,cpu,time,host,mstate,anctaskid
basic    taskid,projid,cpu,time
flow:
extended
saddr,daddr,sport,dport,proto,dsfield,nbytes,npkts,action,ctime,lseen,projid,uid
basic    saddr,daddr,sport,dport,proto,nbytes,npkts,action
```

▼ How to Deactivate Process, Task, and Flow Accounting

To deactivate process, task, and flow accounting, turn off each of them individually.

1. Become superuser.

2. Turn off process accounting.

```
# acctadm -x process
```

3. Turn off task accounting.

```
# acctadm -x task
```

4. Turn off flow accounting.

```
# acctadm -x flow
```

5. Verify that task accounting, process accounting, and flow accounting have been turned off.

```
# acctadm
    Task accounting: inactive
    Task accounting file: none
    Tracked task resources: extended
    Untracked task resources: none
    Process accounting: inactive
    Process accounting file: none
    Tracked process resources: extended
    Untracked process resources: host,mstate
    Flow accounting: inactive
    Flow accounting file: none
    Tracked flow resources: extended
    Untracked flow resources: none
```

Composed September 5, 2002

Resource Controls

After you determine the resource consumption of workloads on your system as described in Chapter 7, you can place bounds on resource usage. Bounds prevent workloads from over-consuming resources. The resource controls facility, which extends the UNIX resource limit concept, is the constraint mechanism that is used for this purpose.

Overview

UNIX systems have traditionally provided a resource limit facility (*rlimit*). The *rlimit* facility allows administrators to set one or more numerical limits on the amount of resources a process can consume. These limits include per-process CPU time used, per-process core file size, and per-process maximum heap size. Heap size is the amount of memory that is allocated for the process data segment.

In the Solaris operating environment, the concept of a per-process resource limit has been extended to the task and project entities described in Chapter 6. The extended limits can be observed on a system-wide basis. These enhancements are provided by the *resource controls* (*rctl*s) facility.

The resource controls facility provides compatibility interfaces for the resource limits facility. Existing applications that use resource limits continue to run unchanged. These applications can be observed in the same way as applications that are modified to take advantage of the resource controls facility.

Resource controls provide a mechanism for constraint on system resources. Processes, tasks, and projects can be prevented from consuming amounts of specified system resources. This mechanism leads to a more manageable system by preventing over-consumption of resources.

Constraint mechanisms can be used to support capacity-planning processes. An encountered constraint can provide information about application resource needs without necessarily denying the resource to the application.

Resource controls can also serve as a simple attribute mechanism for resource management facilities. For example, the number of CPU shares made available to a project in the fair share scheduler (FSS) scheduling class is defined by the `project.cpu-shares` resource control. Because a project is assigned a fixed number of shares by the control, the various actions associated with exceeding a control are not relevant. In this context, the current value for the `project.cpu-shares` control is considered an attribute on the specified project.

Administering Resource Controls

The resource controls facility is configured through the `project` database (see Chapter 6). Resource control attributes are set in the final field of the `project` database entry. The values associated with each resource control are enclosed in parentheses, and appear as plain text separated by commas. The values in parentheses constitute an “action clause.” Each action clause is composed of a privilege level, a threshold value, and an action that is associated with the particular threshold. Each resource control can have multiple action clauses, which are also separated by commas. The following entry defines a per-process address-space limit and a per-task lightweight process limit on a project entity.

```
development:101:Developers:::task.max-lwps=(privileged,10,deny);
process.max-address-space=(privileged,209715200,deny)
```

The `rctladm(1M)` command allows you to make runtime interrogations of and modifications to the resource controls facility, with global scope. The `prctl(1)` command allows you to make runtime interrogations of and modifications to the resource controls facility, with local scope.

Available Resource Controls

A list of the standard resource controls that are available in this release is shown in the following table.

The table describes the resource that is constrained by each control. The table also identifies the default units that are used by the `project` database for that resource. The default units are of two types:

- Quantities represent a limited amount.
- Indexes represent a maximum valid identifier.

Thus, `project.cpu-shares` specifies the number of shares to which the project is entitled. `process.max-file-descriptor` specifies the highest file number that can be assigned to a process by the `open(2)` system call.

TABLE 8-1 Standard Resource Controls

Control Name	Description	Default Unit
<code>project.cpu-shares</code>	The number of CPU shares that are granted to this project for use with <code>FSS(7)</code>	Quantity (shares)
<code>task.max-cpu-time</code>	Maximum CPU time that is available to this task's processes	Time (seconds)
<code>task.max-lwps</code>	Maximum number of LWPs simultaneously available to this task's processes	Quantity (LWPs)
<code>process.max-cpu-time</code>	Maximum CPU time that is available to this process	Time (seconds)
<code>process.max-file-descriptor</code>	Maximum file descriptor index that is available to this process	Index (maximum file descriptor)
<code>process.max-file-size</code>	Maximum file offset that is available for writing by this process	Size (bytes)
<code>process.max-core-size</code>	Maximum size of a core file that is created by this process	Size (bytes)
<code>process.max-data-size</code>	Maximum heap memory that is available to this process	Size (bytes)
<code>process.max-stack-size</code>	Maximum stack memory segment that is available to this process	Size (bytes)
<code>process.max-address-space</code>	Maximum amount of address space, as summed over segment sizes, available to this process	Size (bytes)

Resource Control Values and Privilege Levels

A threshold value on a resource control constitutes an enforcement point where local actions can be triggered or global actions, such as logging, can occur.

Each threshold value must be associated with a privilege level of one of the following three types.

- Basic, which can be modified by the owner of the calling process
- Privileged, which can be modified only by privileged (superuser) callers
- System, which is fixed for the duration of the operating system instance

A resource control is guaranteed to have one system value, which is defined by the system, or resource provider. The system value represents how much of the resource the current implementation of the operating system is capable of providing.

Any number of privileged values can be defined, and only one basic value is allowed. Operations that are performed without specifying a privilege value are assigned a basic privilege by default.

The privilege level for a resource control value is defined in the privilege field of the resource control block as `RCTL_BASIC`, `RCTL_PRIVILEGED`, or `RCTL_SYSTEM`. See `getrctl(2)` for more information. You can use the `prctl` command to modify values that are associated with basic and privileged levels.

Actions on Resource Control Values

For each threshold value that is placed on a resource control, you can associate one or more actions.

- You can choose to deny the resource requests for an amount that is greater than the threshold.
- You can choose to send a signal to the violating or observing process if the threshold value is reached.

Due to implementation restrictions, the global properties of each control can restrict the set of available actions that can be set on the threshold value. A list of available signal actions is presented in the following table. For additional information on signals, see `signal(3HEAD)`.

TABLE 8-2 Signals Available to Resource Control Values

Signal	Notes
SIGABRT	
SIGHUP	
SIGTERM	
SIGKILL	
SIGSTOP	
SIGXRES	

TABLE 8-2 Signals Available to Resource Control Values *(Continued)*

Signal	Notes
SIGXFSZ	Available only to resource controls with the RCTL_GLOBAL_FILE_SIZE property (<code>process.max-file-size</code>). See <code>rctlblk_set_value(3C)</code> for more information.
SIGXCPU	Available only to resource controls with the RCTL_GLOBAL_CPU_TIME property (<code>process.max-cpu-time</code>). See <code>rctlblk_set_value(3C)</code> for more information.

Resource Control Flags and Properties

Each resource control on the system has a certain set of associated properties. This set of properties is defined as a set of global flags, which are associated with all controlled instances of that resource. Global flags cannot be modified, but the flags can be retrieved by using either `rctladm` or the `getrctl` system call.

Local flags define the default behavior and configuration for a specific threshold value of that resource control on a specific process or process collective. The local flags for one threshold value do not affect the behavior of other defined threshold values for the same resource control. However, the global flags affect the behavior for every value associated with a particular control. Local flags can be modified, within the constraints supplied by their corresponding global flags, by the `prctl` command or the `setrctl` system call (see `setrctl(2)`).

For the complete list of local flags, global flags, and their definitions, see `rctlblk_set_value(3C)`.

To determine system behavior when a threshold value for a particular resource control is reached, use `rctladm` to display the global flags for the resource control. For example, to display the values for `process.max-cpu-time`, type the following:

```
$ rctladm process.max-cpu-time
process.max-cpu-time  syslog=off  [ lowerable no-deny cpu-time inf ]
```

The global flags indicate the following.

lowerable	Superuser privileges are not required to lower the privileged values for this control.
no-deny	Even when threshold values are exceeded, access to the resource is never denied.
cpu-time	SIGXCPU is available to be sent when threshold values of this resource are reached.
inf	Any value with RCTL_LOCAL_MAXIMAL defined actually represents an infinite quantity, and the value is never enforced.

Use `prctl` to display local values and actions for the resource control.

```
$ prctl -n process.max-cpu-time $$
353939: -ksh
process.max-cpu-time    [ lowerable no-deny cpu-time inf ]
18446744073709551615 privileged signal=XCPU      [ max ]
18446744073709551615 system      deny           [ max ]
```

The max (RCTL_LOCAL_MAXIMAL) flag is set for both threshold values, and the inf (RCTL_GLOBAL_INFINITE) flag is defined for this resource control. Hence, as configured, both threshold quantities represent infinite values and they will never be exceeded.

Resource Control Enforcement

More than one resource control can exist on a resource. A resource control can exist at each containment level in the process model. If resource controls are active on the same resource at different container levels, the smallest container's control is enforced first. Thus, action is taken on `process.max-cpu-time` before `task.max-cpu-time` if both controls are encountered simultaneously.




FIGURE 8-1 Process Collectives, Container Relationships, and Their Resource Control Sets

Global Monitoring of Resource Control Events

Often, the resource consumption of processes is unknown. To get more information, try using the global resource control actions that are available with `rctladm(1M)`. Use `rctladm` to establish a `syslog` action on a resource control. Then, if any entity managed by that resource control encounters a threshold value, a system message is logged at the configured logging level.

Configuration

Each resource control listed in Table 8–1 can be assigned to a project at login or when `newtask(1)` or the other project-aware launchers `at(1)`, `batch` (see `at(1)`), or `cron(1M)` are invoked. Each command that is initiated is launched in a separate task with the invoking user's default project.

Updates to entries in the `project` database, whether to the `/etc/project` file or to a representation of the database in a network name service, are not applied to currently active projects. The updates are applied when a new task joins the project through `login(1)` or `newtask`.

Temporarily Updating Resource Control Values on a Running System

Values changed in the `project` database only become effective for new tasks that are started in a project. However, you can use the `rctladm` and `prctl` commands to update resource controls on a running system.

Updating Logging Status

The `rctladm` command affects the global logging state of each resource control on a system-wide basis. This command can be used to view the global state and to set up the level of `syslog` logging when controls are exceeded.

Updating Resource Controls

You can view and temporarily alter resource control values and actions on a per-process, per-task, or per-project basis by using `prctl`. A project, task, or process ID is given as input, and the command operates on the resource control at the level where it is defined.

Any modifications to values and actions take effect immediately. However, these modifications apply to the current session only. The changes are not recorded in the project database. If the system is restarted, the modifications are lost. Permanent changes to resource controls must be made in the project database.

All resource control settings that can be modified in the project database can also be modified with the `prctl` command. Both basic and privileged values can be added or be deleted and their actions can be modified. By default, the basic type is assumed for all set operations, but processes and users with superuser privileges can also modify privileged resource controls. System resource controls cannot be altered.

Using Resource Controls

▼ How to Set the Maximum Number of LWPs for Each Task in a Project

Type this entry in the `/etc/project` database to set the maximum number of LWPs in each task in project *x-files* to 3.

```
x-files:100::root::task.max-lwps=(privileged,3,deny)
```

When superuser creates a new task in project *x-files* by joining it with `newtask`, superuser will not be able to create more than three LWPs while running in this task. This is shown in the following annotated sample session.

```
# newtask -p x-files csh

# prctl -n task.max-lwps $$
688: csh
task.max-lwps
                                3 privileged deny
                                2147483647 system    deny

# id -p
uid=0(root) gid=1(other) projid=100(x-files)

# ps -o project,taskid -p $$
PROJECT TASKID
x-files    236

# csh          /* creates second LWP */

# csh          /* creates third LWP */
```



```
# csh          /* cannot create more LWPs */
Vfork failed

#
```

▼ How to Set Multiple Controls on a Project

The `/etc/project` file can contain settings for multiple resource controls for each project as well as multiple threshold values for each control. Threshold values are defined in action clauses, which are comma-separated for multiple values.

The following line in the file sets a `basic` control with no action on the maximum LWPs per task for project *x-files*. The line also sets a privileged deny control on the maximum LWPs per task. This control causes any LWP creation that exceeds the maximum to fail, as shown in the previous example. Finally, the maximum file descriptors per process are limited at the `basic` level, which forces failure of any open call that exceeds the maximum.

```
x-files:101::root::task.max-lwps=(basic,10,none),(privileged,500,deny);
      process.max-file-descriptor=(basic,128,deny)
```

▼ How to Use `prctl`

As superuser, type `prctl` to display the maximum file descriptor for the current shell that is running:

```
# prctl -n process.max-file-descriptor $$
8437:  sh
process.max-file-descriptor      [ lowerable deny ]
                                256 basic          deny
                                65536 privileged deny
                                2147483647 system   deny
```

Use the `prctl` command to temporarily add a new privileged value to deny the use of more than three LWPs per task for the *x-files* project. The result is identical to the result in “How to Set the Maximum Number of LWPs for Each Task in a Project” on page 96, as shown in the following annotated sample session:

```
# newtask -p x-files

# id -p
uid=0(root) gid=1(other) projid=101(x-files)

# prctl -n task.max-lwps -t privileged -v 3 -e deny -i project x-files

# prctl -n task.max-lwps -i project x-files
670:  sh
```

```
task.max-lwps
           3 privileged deny
2147483647 system      deny
```

You can also use `prctl -r` to change the lowest value of a resource control.

```
# prctl -n process.max-file-descriptor -r -v 128 $$
```

▼ How to Use `rctladm`

You can use `rctladm` to enable the global `syslog` attribute of a resource control. When the control is exceeded, notification is logged at the specified `syslog` level. Type the following:

```
# rctladm -e syslog process.max-file-descriptor
```

Capacity Warnings

A global action on a resource control enables you to receive notice of any entity that is tripping over a resource control value.

For example, assume you want to determine whether a web server possesses sufficient CPUs for its typical workload. You could analyze `sar(1)` data for idle CPU time and load average. You could also examine extended accounting data to determine the number of simultaneous processes that are running for the web server process.

However, an easier approach is to place the web server in a task. You can then set a global action, using `syslog`, to notify you whenever a task exceeds a scheduled number of LWPs appropriate for the machine's capabilities.

▼ How to Determine Whether a Web Server Is Allocated Enough CPU Capacity

1. Use the `prctl` command to place a privileged (superuser-owned) resource control on the tasks that contain an `httpd` process. Limit each task's total number of LWPs to 40, and disable all local actions.

```
# prctl -n task.max-lwps -v 40 -t privileged -d all 'pgrep httpd'
```

2. Enable a system log global action on the `task.max-lwps` resource control.

Composed September 5, 2002

```
# rctladm -e syslog task.max-lwps
```

3. Observe whether the workload trips the resource control.

If it does, you will see /var/adm/messages such as:

```
Jan  8 10:15:15 testmachine unix: [ID 859581 kern.notice]  
NOTICE: privileged rctl task.max-lwps exceeded by task 19
```

Composed September 5, 2002

Fair Share Scheduler

An analysis of workload data can indicate that a particular workload or group of workloads is monopolizing CPU resources. If these workloads are not violating resource constraints on CPU usage, you can modify the allocation policy for CPU time on the system. The fair share scheduling class described in this chapter enables you to allocate CPU time based on shares instead of the priority scheme of the timesharing (TS) scheduling class.

Overview

A fundamental job of the operating system is to arbitrate which processes get access to the system's resources. The process scheduler, which is also called the dispatcher, is the portion of the kernel that controls allocation of the CPU to processes. The scheduler supports the concept of scheduling classes. Each class defines a scheduling policy that is used to schedule processes within the class. The default scheduler in the Solaris operating environment, the TS scheduler, tries to give every process relatively equal access to the available CPUs. However, you might want to specify that certain processes be given more resources than others.

You can use the *fair share scheduler* (FSS) to control the allocation of available CPU resources among workloads, based on their importance. This importance is expressed by the number of *shares* of CPU resources that you assign to each workload.

You give each project CPU shares to control the project's entitlement to CPU resources. The FSS guarantees a fair dispersion of CPU resources among projects that is based on allocated shares, independent of the number of processes that are attached to a project. The FSS achieves fairness by reducing a project's entitlement for heavy CPU usage and increasing its entitlement for light usage, in accordance with other projects.

The FSS consists of a kernel scheduling class module and class-specific versions of the `dispadm(1M)` and `priocntl(1)` commands. Project shares used by the FSS are specified through the `project.cpu-shares` property in the `project(4)` database.

CPU Share Definition

The term “share” is used to define a portion of the system’s CPU resources that is allocated to a project. If you assign a greater number of CPU shares to a project, relative to other projects, the project receives more CPU resources from the fair share scheduler.

CPU shares are not equivalent to percentages of CPU resources. Shares are used to define the relative importance of workloads in relation to other workloads. When you assign CPU shares to a project, your primary concern is not the number of shares the project has. Knowing how many shares the project has in comparison with other projects is more important. You must also take into account how many of those other projects will be competing with it for CPU resources.

Note – Processes in projects with zero shares always run at the lowest system priority (0). These processes only run when projects with nonzero shares are not using CPU resources.

CPU Shares and Process State

In the Solaris operating environment, a project workload usually consists of more than one process. From the fair share scheduler perspective, each project workload can be in either an *idle* state or an *active* state. A project is considered idle if none of its processes are using any CPU resources. This usually means that such processes are either *sleeping* (waiting for I/O completion) or stopped. A project is considered active if at least one of its processes is using CPU resources. The sum of shares of all active projects is used in calculating the portion of CPU resources to be assigned to projects.

The following formula shows how the FSS scheduler calculates per-project allocation of CPU resources.

$$\text{allocation}_{\text{project } i} = \frac{\text{shares}_{\text{project } i}}{\sum_{j=1 \dots n} (\text{shares}_{\text{project } j})}$$

j is the index among all active projects

FIGURE 9-1 FSS Scheduler Share Calculation

When more projects become active, each project's CPU allocation is reduced, but the proportion between the allocations of different projects does not change.

CPU Share Versus Utilization

Share allocation is not the same as utilization. A project that is allocated 50 percent of the CPU resources might average only a 20 percent CPU use. Moreover, shares serve to limit CPU usage only when there is competition from other projects. Regardless of how low a project's allocation is, it always receives 100 percent of the processing power if it is running alone on the system. Available CPU cycles are never wasted. They are distributed between projects.

The allocation of a small share to a busy workload might slow its performance. However, the workload is not prevented from completing its work if the system is not overloaded.

CPU Share Examples

Assume you have a system with two CPUs running two parallel CPU-bound workloads called A and B , respectively. Each workload is running as a separate project. The projects have been configured so that project A is assigned S_A shares, and project B is assigned S_B shares.


On average, under the traditional TS scheduler, each of the workloads that is running on the system would be given the same amount of CPU resources. Each workload would get 50 percent of the system's capacity.

When run under the control of the FSS scheduler with $S_A=S_B$, these projects are also given approximately the same amounts of CPU resources. However, if the projects are given different numbers of shares, their CPU resource allocations are different.

The next three examples illustrate how shares work in different configurations. These examples show that shares are only mathematically accurate for representing the usage if demand meets or exceeds available resources.

Example 1: Two CPU-Bound Processes in Each Project

If *A* and *B* each have two CPU-bound processes, and $S_A = 1$ and $S_B = 3$, then the total number of shares is $1 + 3 = 4$. In this configuration, given sufficient CPU demand, projects *A* and *B* are allocated 25 percent and 75 percent of CPU resources, respectively.



Example 2: No Competition Between Projects

If *A* and *B* have only *one* CPU-bound process each, and $S_A = 1$ and $S_B = 100$, then the total number of shares is 101. Each project cannot use more than one CPU because each project has only one running process. Because no competition exists between projects for CPU resources in this configuration, projects *A* and *B* are each allocated 50 percent of all CPU resources. In this configuration, CPU share values are irrelevant. The projects' allocations would be the same (50/50), even if both projects were assigned zero shares.

50%	50%
(1st CPU)	(2nd CPU)
<i>Project A</i> (1 share)	<i>Project B</i> (100 shares)

Example 3: One Project Unable to Run

If *A* and *B* have two CPU-bound processes each, and project *A* is given 1 share and project *B* is given 0 shares, then project *B* is not allocated any CPU resources and project *A* is allocated all CPU resources. Processes in *B* always run at system priority 0, so they will never be able to run because processes in project *A* always have higher priorities.

100%	0%
<i>Project A</i> (1 share)	<i>Project B</i> (0 shares)

FSS Configuration

Projects and Users

Projects are the workload containers in the FSS scheduler. Groups of users who are assigned to a project are treated as single controllable blocks. Note that you can create a project with its own number of shares for an individual user.

Users can be members of multiple projects that have different numbers of shares assigned. By moving processes from one project to another project, processes can be assigned CPU resources in varying amounts.

For more information on the `project(4)` database and name services, see “`project Database`” on page 70.

CPU Shares Configuration

The configuration of CPU shares is managed by the name service as a property of the `project` database.

When the first task (or process) associated with a project is created through the `setproject(3PROJECT)` library function, the number of CPU shares defined as resource control `project.cpu-shares` in the `project` database is passed to the kernel. A project that does not have the `project.cpu-shares` resource control defined is assigned one share.

In the following example, this entry in the `/etc/project` file sets the number of shares for project *x-files* to 5:

```
x-files:100:::project.cpu-shares=(privileged,5,none)
```

If you alter the number of CPU shares allocated to a project in the database when processes are already running, the number of shares for that project will not be modified at that point. The project must be restarted for the change to become effective.

If you want to temporarily change the number of shares assigned to a project without altering the project's attributes in the `project` database, use `prctl(1)`. For example, to change the value of project *x-files*'s `project.cpu-shares` resource control to 3 while processes associated with that project are running, type the following:

```
# prctl -r -n project.cpu-shares -v 3 -i project x-files
```

`-r` Replaces the current value for the named resource control.

- n *name* Specifies the name of the resource control.
- v *val* Specifies the value for the resource control.
- i *idtype* Specifies the ID type of the next argument.
- x-files* Specifies the object of the change. In this instance, project *x-files* is the object.

Project *system* with project ID 0 includes all system daemons that are started by the boot-time initialization scripts. *system* can be viewed as a project with an unlimited number of shares. This means that *system* is always scheduled first, regardless of how many shares have been given to other projects. If you do not want the *system* project to have unlimited shares, you can specify a number of shares for this project in the *project* database.

As stated previously, processes that belong to projects with zero shares are always given zero system priority. Projects with one or more shares are running with priorities one and higher. Thus, projects with zero shares are only scheduled when CPU resources are available that are not requested by a nonzero share project.

The maximum number of shares that can be assigned to one project is 65535.

FSS and Processor Sets

The FSS can be used in conjunction with processor sets to provide more fine-grained controls over allocations of CPU resources among projects that run on each processor set than would be available with processor sets alone. The FSS scheduler treats processor sets as entirely independent partitions, with each processor set controlled independently with respect to CPU allocations.

The CPU allocations of projects running in one processor set are not affected by the CPU shares or activity of projects running in another processor set because the projects are not competing for the same resources. Projects only compete with each other if they are running within the same processor set.

The number of shares that is allocated to a project is system wide. Regardless of which processor set it is running on, each portion of a project is given the same amount of shares.

When processor sets are used, project CPU allocations are calculated for active projects that run within each processor set, as shown in the following figure.

$$\text{allocation}_{\text{project } x}^i = \frac{\text{shares}_{\text{project } x}^i}{\sum_{j=1 \dots n} (\text{shares}_{\text{project } j}^i)}$$

j is the index among all active projects that run on processor set X

FIGURE 9-2 FSS Scheduler Share Calculation With Processor Sets

Project partitions that run on different processor sets might have different CPU allocations. The CPU allocation for each project partition in a processor set depends only on the allocations of other projects that run on the same processor set.

The performance and availability of applications that run within the boundaries of their processor sets are not affected by the introduction of new processor sets. The applications are also not affected by changes that are made to the share allocations of projects that run on other processor sets.

Empty processor sets (sets without processors in them) or processor sets without processes bound to them do not have any impact on the FSS scheduler behavior.

FSS and Processor Sets Examples

Assume that a server with eight CPUs is running several CPU-bound applications in projects A , B , and C . Project A is allocated one share, project B is allocated two shares, and project C is allocated three shares.

Project A is running only on processor set 1. Project B is running on processor sets 1 and 2. Project C is running on processor sets 1, 2, and 3. Assume that each project has enough processes to utilize all available CPU power. Thus, there is always competition for CPU resources on each processor set.

Composed September 5, 2002

Project A 16.66% (1/6)	Project B 40% (2/5)	Project C 100% (3/3)
Project B 33.33% (2/6)		
Project C 50% (3/6)	Project C 60% (3/5)	
Processor Set #1 2 CPUs 25% of the system	Processor Set #2 4 CPUs 50% of the system	Processor Set #3 2 CPUs 25% of the system

The total system-wide project CPU allocations on such a system are as follows:

$$\begin{aligned}
 \text{Project A} & \quad 4\% = (1/6 \times 2/8)_{\text{pset1}} \\
 \text{Project B} & \quad 28\% = (2/6 \times 2/8)_{\text{pset1}} + (2/5 \times 4/8)_{\text{pset2}} \\
 \text{Project C} & \quad 67\% = (3/6 \times 2/8)_{\text{pset1}} + (3/5 \times 4/8)_{\text{pset2}} + (3/3 \times 2/8)_{\text{pset3}}
 \end{aligned}$$

These percentages do not match the corresponding amounts of CPU shares that are given to projects. However, within each processor set, the per-project CPU allocation ratios are proportional to their respective shares.

On the same system *without* processor sets, the distribution of CPU resources would be different:

$$\begin{aligned}
 \text{Project A} & \quad 16.66\% = (1/6) \\
 \text{Project B} & \quad 33.33\% = (2/6) \\
 \text{Project C} & \quad 50\% = (3/6)
 \end{aligned}$$

Combining FSS With Other Scheduling Classes

By default, the FSS scheduling class uses the same range of priorities (0 to 59) as the timesharing (TS), interactive (IA), and fixed priority (FX) scheduling classes. Therefore, you should avoid having processes from these scheduling classes share *the same* processor set. A mix of processes in the FSS, TS, IA, and FX classes could result in unexpected scheduling behavior.

With the use of processor sets, you can mix TS, IA, and FX with FSS in one system. However, all the processes that run on each processor set must be in *one* scheduling class, so they do not compete for the same CPUs. The FX scheduler in particular should not be used in conjunction with the FSS scheduling class unless processor sets are used. This action prevents applications in the FX class from using priorities high enough to starve applications in the FSS class.

You can mix processes in the TS and IA classes in the same processor set, or on the same system without processor sets.

The Solaris operating environment also offers a real-time (RT) scheduler to users with superuser privileges. By default, the RT scheduling class uses system priorities in a different range (usually from 100 to 159) than FSS. Because RT and FSS are using disjoint ranges of priorities, FSS can coexist with the RT scheduling class within the same processor set. However, the FSS scheduling class does not have any control over processes that run in the RT class.

For example, on a four-processor system, a single-threaded RT process can consume one entire processor if the process is CPU bound. If the system also runs FSS, regular user processes compete for the three remaining CPUs that are not being used by the RT process. Note that the RT process might not use the CPU continuously. When the RT process is idle, FSS utilizes all four processors.

You can type the following command to find out which scheduling classes the processor sets are running in and ensure that each processor set is configured to run either TS, IA, FX, or FSS processes.

```
$ ps -ef -o pset,class | grep -v CLS | sort | uniq
1 FSS
1 SYS
2 TS
2 RT
3 FX
```

To set the default scheduler for the system, see “FSS Configuration Examples” on page 111 and `dispadm(1M)`. To move running processes into a different scheduling class, see “FSS Configuration Examples” on page 111 and `pricntl(1)`.

Monitoring the FSS

You can use `prstat(1M)` to monitor CPU usage by active projects.

You can use the extended accounting data for tasks to obtain per-project statistics on the amount of CPU resources that are consumed over longer periods. See Chapter 7 for more information.

▼ How to Monitor System CPU Usage by Projects

To monitor the CPU usage of projects that run on the system, type the following:

```
% prstat -J
```

▼ How to Monitor CPU Usage by Projects in Processor Sets

To monitor the CPU usage of projects on a list of processor sets, type the following:

```
% prstat -J -C pset-list
```

pset-list List of processor set IDs that are separated by commas

FSS Configuration Examples

As with other scheduling classes in the Solaris environment, commands to set the scheduler class, configure the scheduler's tunable parameters, and configure the properties of individual processes can be used with FSS.

▼ How to Set the Scheduler Class

Use the `dispadm` command to set FSS as the default scheduler for the system.

```
# dispadm -d FSS
```

This change takes effect on the next reboot. After reboot, every process on the system runs in the FSS scheduling class.

▼ How to Manually Move Processes From the TS Into the FSS Class

You can manually move processes from the TS scheduling class into the FSS scheduling class without changing the default scheduling class and rebooting.

1. **Become superuser.**

2. **Move the `init` process (pid 1) into the FSS scheduling class.**

```
# priocntl -s -c FSS -i pid 1
```

3. **Move all processes from the TS scheduling class into the FSS scheduling class.**

```
# priocntl -s -c FSS -i class TS
```

All processes again run in the TS scheduling class after reboot.

▼ How to Manually Move Processes From all User Classes Into the FSS Class

You might be using a default class other than TS. For example, your system might be running a window environment that uses the IA class by default. You can manually move all processes into the FSS scheduling class without changing the default scheduling class and rebooting.

1. **Become superuser.**

2. **Move the `init` process (pid 1) into the FSS scheduling class.**

```
# priocntl -s -c FSS -i pid 1
```

3. **Move all processes from their current scheduling classes into the FSS scheduling class.**

```
# priocntl -s -c FSS -i all
```

All processes again run in the default scheduling class after reboot.

▼ How to Move a Project's Processes Into the FSS Class

You can manually move processes in a particular project from their current scheduling class to the FSS scheduling class.

1. Become superuser.
2. Move processes that run in project ID 10 to the FSS scheduling class.

```
# priocntl -s -c FSS -i projid 10
```

The project's processes again run in the default scheduling class after reboot.

▼ How to Tune Scheduler Parameters

You can use the `dispadmin` command to examine and tune the FSS scheduler's time quantum value. *Time quantum* is the amount of time that a thread is allowed to run before it must relinquish the processor. To display the current time quantum for the FSS scheduler, type the following:

```
$ dispadmin -c FSS -g
#
# Fair Share Scheduler Configuration
#
RES=1000
#
# Time Quantum
#
QUANTUM=110
```

When you use the `-g` option, you can also use the `-r` option to specify the resolution that is used for printing time quantum values. If no resolution is specified, time quantum values are displayed in milliseconds by default. Type the following:

```
$ dispadmin -c FSS -g -r 100
#
# Fair Share Scheduler Configuration
#
RES=100
#
# Time Quantum
#
QUANTUM=11
```

To set scheduling parameters for the FSS scheduling class, use `dispadmin -s`. The values in *file* must be in the format output by the `-g` option. These values overwrite the current values in the kernel. Type the following:

```
$ dispadmin -c FSS -s file
```

Composed September 5, 2002

References

For more information on how to use the FSS scheduler, see `priocntl(1)`, `ps(1)`, `dispadm(1M)`, and `FSS(7)`.

Resource Pools

This chapter discusses resource pools, which are used for partitioning machine resources. Resource pools enable you to separate workloads so that workload consumption of certain resources does not overlap. This resource reservation helps to achieve predictable performance on systems with mixed workloads.

Overview

Resource pools provide a persistent configuration mechanism for processor set configuration and, optionally, scheduling class assignment.




FIGURE 10-1 Resource Pool Framework

By grouping multiple partitions, pools provide a handle to associate with labeled workloads. Each project entry in the `/etc/project` database can have a pool associated with it. New work that is started on a project is bound to the appropriate pool.

The pools mechanism is primarily for use on large machines of more than four CPUs. However, small machines can still benefit from this functionality. On small machines, you can create pools that share noncritical resource partitions. The pools are separated only on the basis of critical resources.

When to Use Pools

Resource pools offer a versatile mechanism that can be applied to many administrative scenarios, as described in the following sections.

Batch Compute Server

Use pools functionality to split a server into two pools.

One pool is used for login sessions and interactive work by timesharing users. The other pool is used for jobs that are submitted through the batch system.

Application or Database Server

Partition the resources for interactive applications in accordance with the applications' requirements.

Turning on Applications in Phases

Set user expectations.

You might initially deploy a machine that is running only a fraction of the services that the machine is ultimately expected to deliver. User difficulties can occur if reservation-based resource management mechanisms are not established when the machine comes online.

For example, the fair share scheduler optimizes CPU utilization. The response times for a machine that is running only one application can be misleadingly fast when compared to the response times users see with multiple applications loaded. By using separate pools for each application, you can ensure that a ceiling on the number of CPUs available to each application is in place before all applications are deployed.

Complex Timesharing Server

Partition a server that supports large user populations.

Server partitioning provides an isolation mechanism that leads to a more predictable per-user response.

By dividing users into groups that bind to separate pools, and using the fair share scheduling (FSS) facility, you can tune CPU allocations to favor sets of users that have priority. This assignment can be based on user role, accounting chargeback, and so forth.

Workloads That Change Seasonally

Use resource pools to adjust to changing demand.

Your site might experience predictable shifts in workload demand over long periods of time, such as monthly, quarterly, or annual cycles. If your site experiences these shifts, you can alternate between multiple pools configurations by invoking `pooladm` from a `cron(1M)` job.

Real-Time Applications

Create a real-time pool by using the RT scheduler and designated processor resources.

Administering Pools

The commands that are shown in the following table provide the primary administrative interface to the pools facility.

Command	Description
<code>pooladm(1M)</code>	Activates a particular configuration or deactivates the current configuration. If run without options, <code>pooladm</code> prints out the current running pools configuration.
<code>poolbind(1M)</code>	Enables the manual binding of projects, tasks, and processes to a pool.
<code>poolcfg(1M)</code>	Creates and modifies pools configuration files. If run with the <code>info</code> subcommand argument to the <code>-c</code> option, <code>poolcfg</code> displays the current configuration.

A library API is provided by `libpool(3LIB)`. The library can be used by programs to manipulate pool configurations.

Pools Framework

The resource pools framework stores its view of the machine in a private configuration file. (The location of the file is private to the implementation of the pools framework.) This configuration file represents the pools framework's view of the machine. The file also contains information about configured pools and the organization of partitionable resources. Each pool can contain the following:

- A reference to a processor set or a CPU resource partition
- A property that specifies the pool's default scheduling class

Implementing Pools on a System

Pools can be implemented on a system by using one of these methods.

1. When the Solaris software boots, an `init` script checks to see if the `/etc/pooladm.conf` file exists. If this file is found, then `pooladm` is invoked to make this configuration the active pools configuration. The system creates a private configuration file to reflect the organization that is requested in `/etc/pooladm.conf`, and the machine's resources are partitioned accordingly.
2. When the Solaris environment is running, a pools configuration can either be activated if it is not already present, or modified by using the `pooladm` command. By default, `pooladm` operates on `/etc/pooladm.conf`. However, you can optionally specify an alternate location and file name, and use this file to update the pools configuration.

Dynamic Reconfiguration Operations and Resource Pools

Dynamic reconfiguration (DR) enables you to reconfigure hardware while the system is running. Because DR affects available resource amounts, the pools facility must be included in these operations. When a DR operation is initiated, the pools framework acts to validate the configuration.

If the DR operation can proceed without causing the current pools configuration to become invalid, then the private configuration file is updated. An invalid configuration is one that cannot be supported by the available resources.

If the DR operation would cause the pools configuration to be invalid, then the operation fails and you are notified by a message to the message log. If you want to force the configuration to completion, you must use the DR force option. The pools configuration is then modified to comply with the new resource configuration.

Creating Pools Configurations

The configuration file contains a description of the pools to be created on the system. The file describes the entities and resource types that can be manipulated.

Type	Description
pset	A processor set resource
pool	Named collection of resource associations
system	The machine-level entity

See `poolcfg(1M)` for more information on elements that be manipulated.

You can create a structured `/etc/pooladm.conf` file in two ways.

- You can use `poolcfg` to discover the resources on the current system and place the results in a configuration file.

This method simplifies file construction. All active resources and components on the system that are capable of being manipulated by the pools facility are recorded. The resources include existing processor set configurations. You can then modify the configuration to rename the processor sets or to create additional pools if necessary.

- You can use `poolcfg` to create a new pools configuration.

Use this method when you develop configurations for other machines or when you create configurations that you want to apply to the current machine at a later time.

Use `poolcfg` or `libpool` to modify the `/etc/pooladm.conf` file. Do not directly edit this file.

▼ How to Create a Configuration by Discovery

Use the `discover` subcommand argument to the `-c` option of `/usr/sbin/poolcfg` to create the pools configuration file. The resulting file, `/etc/pooladm.conf`, contains any existing processor sets.

1. **Become superuser.**
2. **Type the following:**

```
# poolcfg -c discover
```

You can also supply a file name to use instead of the default `/etc/pooladm.conf`. If the file name is supplied, then the `poolcfg` commands are applied to the contents of the named file.

For example, to place a discovered configuration in `/tmp/foo`, do the following:

1. **Become superuser.**
2. **Type the following:**

```
# poolcfg -c discover /tmp/foo
```

▼ How to Create a New Configuration

Use the `create` subcommand argument to the `-c` option of `/usr/sbin/poolcfg` to create a simple configuration file for a system named `tester`. Note that you must quote subcommand arguments that contain white space.

1. **Become superuser.**
2. **Type the following:**
3. **View the contents of the configuration file in readable form.**

```
# poolcfg -c 'create system tester'

system tester
    int system.version 1
    boolean system.bind-default true
    string system.comment
```

▼ How to Modify a Configuration

To enhance your simple configuration, create a processor set named *batch* and a pool named *batch*. Then join them with an association. Note that you must quote subcommand arguments that contain white space.

1. Become superuser.

2. Create processor set `batch`.

```
# poolcfg -c 'create pset batch (uint pset.min = 2; uint pset.max = 10)'
```

3. Create pool `batch`.

```
# poolcfg -c 'create pool batch'
```

4. Join with an association.

```
# poolcfg -c 'associate pool batch (pset batch)'
```

5. Display the edited configuration.

```
# poolcfg -c info
system tester
  int system.version 1
  boolean system.bind-default true
  string system.comment

  pool batch
    boolean pool.default false
    boolean pool.active true
    int pool.importance 1
    string pool.comment
    pset batch

  pset batch
    int pset.sys_id -2
    string pset.units population
    boolean pset.default true
    uint pset.max 10
    uint pset.min 2
    string pset.comment
    boolean pset escapable false
    uint pset.load 0
    uint pset.size 0
```

▼ How to Associate a Pool With a Scheduling Class

You can associate a pool with a scheduling class so that all processes bound to the pool use this scheduler. To do this, set the `pool.scheduler` property to the name of the scheduler class. This example shows how to associate the pool `batch` with the FSS.

1. Become superuser.

2. Modify pool `batch` to be associated with the FSS.

```
# poolcfg -c 'modify pool batch (string pool.scheduler="FSS")'
```

3. Display the edited configuration.

```
# poolcfg -c info
system tester
    int system.version 1
    boolean system.bind-default true
    string system.comment

    pool batch
        boolean pool.default false
        boolean pool.active true
        int pool.importance 1
        string pool.comment
        string pool.scheduler FSS
        pset batch

    pset batch
        int pset.sys_id -2
        string pset.units population
        boolean pset.default true
        uint pset.max 10
        uint pset.min 2
        string pset.comment
        boolean pset.escapable false
        uint pset.load 0
        uint pset.size 0
```

▼ How to Use Command Files With poolcfg

poolcfg -f can take input from a text file that contains poolcfg subcommand arguments to the -c option. This technique is appropriate when you want a set of operations to be performed atomically. When processing multiple commands, the configuration is only updated if all of the commands succeed. For large or complex configurations, this technique can be more useful than per-subcommand invocations.

1. Create the input file.

```
$ cat > poolcmds.txt
create system tester
create pset batch (uint pset.min = 2; uint pset.max = 10)
create pool batch
associate pool batch (pset batch)
```

2. Become superuser.

3. Type the following:

```
# /usr/sbin/poolcfg -f poolcmds.txt
```

Activating and Deactivating Pools Configurations

Use `pooladm(1M)` to make a particular pool configuration active or to remove an active pools configuration.

▼ How to Activate a Pools Configuration

To activate the configuration in the default static configuration file, `/etc/pooladm.conf`, invoke `pooladm` with the `-c` option, “commit configuration.”

1. **Become superuser.**
2. **Type the following:**

```
# /usr/sbin/pooladm -c
```

▼ How to Deactivate a Pools Configuration

To remove the running configuration and all associated resources, such as processor sets, use the `-x` option for “remove configuration.”

1. **Become superuser.**
2. **Type the following:**

```
# /usr/sbin/pooladm -x
```

The `-x` option to `pooladm` removes the dynamic private configuration file as well as all resource configurations that are associated with the dynamic configuration. Thus, the `-x` option provides a mechanism for recovering from a poorly designed pools configuration. All processes share all of the resources on the machine.

Note – Mixing scheduling classes within one processor set can lead to unpredictable results. If you use `pooladm -x` to recover from a bad configuration, you should then use `priocntl(1)` to move running processes into a different scheduling class.

Binding to a Pool

You can bind a running process to a pool in two ways.

- You can use the `poolbind(1M)` command to bind a specific process to a named resource pool.
- You can use the `project.pool` attribute in the `project(4)` database to identify the pool binding for a new login session or a task that is launched through `newtask(1)`.

▼ How to Bind Processes to a Pool

The following procedure manually binds a process (for example, the current shell) to a pool named *ohare*.

1. **Become superuser.**
2. **Type the following:**

```
# poolbind -p ohare $$
```

▼ How to Bind Tasks or Projects to a Pool

To bind tasks or projects to a pool, use `poolbind` with the `-i` option. The following example binds all processes in the *airmiles* project to the *laguardia* pool.

1. **Become superuser.**
2. **Type the following:**

```
# poolbind -i project -p laguardia airmiles
```

▼ How to Use project Attributes to Bind New Processes to a Pool

To automatically bind new processes in a project to a pool, add the `project.pool` attribute to each entry in the `project` database.

For example, assume you have a configuration with two pools that are named *studio* and *backstage*. The `/etc/project` file has the following contents.

```
user.paul:1024:::project.pool=studio
user.george:1024:::project.pool=studio
user.ringo:1024:::project.pool=backstage
```

Composed September 5, 2002

```
passes:1027::paul::project.pool=backstage
```

With this configuration, processes that are started by user `paul` are bound by default to the `studio` pool.

▼ How to Use `project` Attributes to Bind a Process to a Different Pool

Using the previous configuration, user `paul` can modify the pool binding for processes he starts. He can use `newtask` to bind work to the `backstage` pool as well, by launching in the `passes` project.

1. Launch a process in the `passes` project.

```
$ newtask -l -p passes
```

2. Verify the pool binding for the process.

```
$ poolbind -q $$  
process id 6384 : pool 'backstage'
```

Composed September 5, 2002

CHAPTER **11**

Resource Management Configuration Example

This chapter reviews the resource management framework and describes a hypothetical server consolidation project. In this example, five applications are being consolidated onto a single system. The target applications have resource requirements that vary, different user populations, and different architectures.

Configuration to Be Consolidated

Currently, each application exists on a dedicated server that is designed to meet the requirements of the application. The applications and their characteristics are identified in the following table.

Application Description	Characteristics
Application server	Exhibits negative scalability beyond 2 CPUs
Database instance for application server	Heavy transaction processing
Application server in test and development environment	GUI-based, with untested code execution
Transaction processing server	Primary concern is response time
Standalone database instance	Processes a large number of transactions and serves multiple time zones

Consolidation Configuration

The following configuration is used to consolidate the applications onto a single system.

- The application server has a two-CPU processor set.
- The database instance for the application server and the standalone database instance are consolidated onto a single processor set of at least four CPUs. The standalone database instance is guaranteed 75 percent of that resource.
- The test and development application server requires the IA scheduling class to ensure UI responsiveness. Memory limitations are imposed to lessen the effects of bad code builds.
- The transaction processing server is assigned a dedicated processor set of at least two CPUs, to minimize response latency.

Creating the Configuration

Edit the `project` database file. Add entries to implement the required resource controls and to map users to resource pools, and then view the file.

```
# cat /etc/project
.
.
.
user.app_server:2001:Production Application Server::project.pool=appserver_pool
user.app_db:2002:App Server DB::project.pool=db_pool,project.cpu-shares(privileged,1,deny)
development:2003:Test and development::staff:project.pool=dev_pool,
    process.max-address-space=(privileged,536870912,deny)
user.tp_engine:2004:Transaction Engine::project.pool=tp_pool
user.geo_db:2005:EDI DB::project.pool=db_pool,project.cpu-shares=(privileged,3,deny)
.
.
.
```

Note – The development team has to execute tasks in the development project because access for this project is based on a user's group ID (GID).

Create an input file named `pool.host`, which will be used to configure the required resource pools. View the file.

Composed September 5, 2002

```
# cat pool.host
```

```
create system host
create pset default_pset (uint pset.min = 1)
create pset dev_pset (uint pset.max = 2)
create pset tp_pset (uint pset.min = 2)
create pset db_pset (uint pset.min = 4; uint pset.max = 6)
create pset app_pset (uint pset.min = 1; uint pset.max = 2)
create pool default_pool (string pool.scheduler="TS"; boolean pool.default = true)
create pool dev_pool (string pool.scheduler="IA")
create pool appserver_pool (string pool.scheduler="TS")
create pool db_pool (string pool.scheduler="FSS")
create pool tp_pool (string pool.scheduler="TS")
associate pool default_pool (pset default_pset)
associate pool dev_pool (pset dev_pset)
associate pool appserver_pool (pset app_pset)
associate pool db_pool (pset db_pset)
associate pool tp_pool (pset tp_pset)
```

Type the following:

```
# poolcfg -f pool.host
```

Make the configuration active.

```
# pooladm -c
```

The framework is now functional on the system.

Viewing the Configuration

To view the framework configuration, type:

```
# pooladm
system host
    int system.version 1
    boolean system.bind-default true
    string system.comment

    pool default_pool
        boolean pool.default true
        boolean pool.active true
        int pool.importance 1
        string pool.comment
        string pool.scheduler TS
        pset default_pset

    pool dev_pool
        boolean pool.default false
        boolean pool.active true
```

Composed September 5, 2002

```
        int pool.importance 1
        string pool.comment
        string pool.scheduler IA
        pset dev_pset

pool appserver_pool
    boolean pool.default false
    boolean pool.active true
    int pool.importance 1
    string pool.comment
    string pool.scheduler TS
    pset app_pset

pool db_pool
    boolean pool.default false
    boolean pool.active true
    int pool.importance 1
    string pool.comment
    string pool.scheduler FSS
    pset db_pset

pool tp_pool
    boolean pool.default false
    boolean pool.active true
    int pool.importance 1
    string pool.comment
    string pool.scheduler TS
    pset tp_pset

pset default_pset
    int pset.sys_id -1
    string pset.units population
    boolean pset.default true
    uint pset.max 4294967295
    uint pset.min 1
    string pset.comment
    boolean pset.escapable false
    uint pset.load 0
    uint pset.size 0

pset dev_pset
    int pset.sys_id 1
    string pset.units population
    boolean pset.default false
    uint pset.max 2
    uint pset.min 0
    string pset.comment
    boolean pset.escapable false
    uint pset.load 0
    uint pset.size 0

pset tp_pset
    int pset.sys_id 2
    string pset.units population
    boolean pset.default false
```

Composed September 5, 2002

```
uint pset.max 4294967295
uint pset.min 2
string pset.comment
boolean pset.escapable false
uint pset.load 0
uint pset.size 0

pset db_pset
  int pset.sys_id 3
  string pset.units population
  boolean pset.default false
  uint pset.max 6
  uint pset.min 4
  string pset.comment
  boolean pset.escapable false
  uint pset.load 0
  uint pset.size 0

pset app_pset
  int pset.sys_id 4
  string pset.units population
  boolean pset.default false
  uint pset.max 2
  uint pset.min 1
  string pset.comment
  boolean pset.escapable false
  uint pset.load 0
  uint pset.size 0
```

A graphic representation of the framework follows.




FIGURE 11-1 Server Consolidation Configuration

Note – In the `db_pool`, the standalone database instance is guaranteed 75 percent of the CPU resource.

CHAPTER **12**

Resource Control Functionality in the Solaris Management Console

This chapter describes the resource control and performance monitoring features in the Solaris Management Console.

You can use the console to monitor system performance and to enter resource control values for projects, tasks, and processes. The console provides a convenient, secure alternative to the command-line interface (CLI) for managing hundreds of configuration parameters that are spread across many systems. Each system is managed individually. The console's graphical interface supports all experience levels.

Using the Console (Task Map)

Task	Description	For Instructions
Use the console	Start the Solaris Management Console in a local environment or in a name service or directory service environment. Note that the performance tool is not available in a name service environment.	"Starting the Solaris Management Console" in <i>System Administration Guide: Basic Administration</i> and "Using the Solaris Management Tools in a Name Service Environment (Task Map)" in <i>System Administration Guide: Basic Administration</i>
Monitor system performance	Access the Performance tool under System Status.	"How to Access the Performance Tool" on page 135
Add resource controls to projects	Access the Resource Controls tab under System Configuration.	"How to Access the Resource Controls Tab" on page 139

Overview

Resource management functionality is a component of the Solaris Management Console. The console is a container for GUI-based administrative tools that are stored in collections called toolboxes. For information on the console and how to use it, see “Working With the Management Console (Tasks)” in *System Administration Guide: Basic Administration*.

When you use the console and its tools, the main source of documentation is the online help system in the console itself. For a description of the documentation available in the online help, see “Solaris Management Console (Overview)” in *System Administration Guide: Basic Administration*.

Management Scope

The term *management scope* refers to the name service environment that you choose to use with the selected management tool. The management scope choices for the resource control and performance tools are the `/etc/project` local file, or NIS.

The management scope that you select during a console session should correspond to the primary name service that is identified in the `/etc/nsswitch.conf` file.

Performance Tool

The Performance tool is used to monitor resource utilization. Resource utilization can be summarized for the system, viewed by project, or viewed for an individual user.




FIGURE 12-1 Performance Tool in the Solaris Management Console

▼ How to Access the Performance Tool

The Performance tool is located under System Status in the Navigation pane. To access the Performance tool, do the following:

1. **Click the System Status control entity in the Navigation pane.**
The control entity is used to expand menu items in the Navigation pane.
2. **Click the Performance control entity.**
3. **Click the System control entity.**
4. **Double-click Summary, Projects, or Users.**
Your choice depends on the usage you want to monitor.

Monitoring by System

Values are shown for the following attributes.

Attribute	Description
Active Processes	Number of processes that are active on the system
Physical Memory Used	Amount of system memory that is in use
Physical Memory Free	Amount of system memory that is available
Swap Used	Amount of system swap space that is in use
Swap Free	Amount of free system swap space
Page Rate	Rate of system paging activity
System Calls	Number of system calls per second
Network Packets	Number of network packets that are transmitted per second
CPU Usage	Percentage of CPU that is currently in use
Load Average	Number of processes in the system run queue which are averaged over the last 1, 5, and 15 minutes

Monitoring by Project or User Name

Values are shown for the following attributes.

Attribute	Short Name	Description
Input Blocks	inblk	Number of blocks read
Blocks Written	oublk	Number of blocks written
Chars Read/Written	ioch	Number of characters read and written
Data Page Fault Sleep Time	dftime	Amount of time spent processing data page faults
Involuntary Context Switches	ictx	Number of involuntary context switches
System Mode Time	stime	Amount of time spent in the kernel mode

Composed September 5, 2002

Attribute	Short Name	Description
Major Page Faults	majfl	Number of major page faults
Messages Received	mrcv	Number of messages received
Messages Sent	msend	Number of messages sent
Minor Page Faults	minf	Number of minor page faults
Num Processes	nprocs	Number of processes owned by the user or the project
Num LWPs	count	Number of lightweight processes
Other Sleep Time	slptime	Sleep time other than tftime, dftime, kftime, and ltime
CPU Time	pctcpu	Percentage of recent CPU time used by the process, the user, or the project
Memory Used	pctmem	Percentage of system memory used by the process, the user, or the project
Heap Size	brksize	Amount of memory allocated for the process data segment
Resident Set Size	rsssize	Current amount of memory claimed by the process
Process Image Size	size	Size of the process image in Kbytes
Signals Received	sigs	Number of signals received
Stopped Time	stoptime	Amount of time spent in the stopped state
Swap Operations	swaps	Number of swap operations in progress
System Calls Made	sysc	Number of system calls made over the last time interval
System Page Fault Sleep Time	kftime	Amount of time spent processing page faults
System Trap Time	ttime	Amount of time spent processing system traps
Text Page Fault Sleep Time	tftime	Amount of time spent processing text page faults

Attribute	Short Name	Description
User Lock Wait Sleep Time	ltime	Amount of time spent waiting for user locks
User Mode Time	utime	Amount of time spent in the user mode
User and System Mode Time	time	The cumulative CPU execution time
Voluntary Context Switches	vctx	Number of voluntary context switches
Wait CPU Time	wtime	Amount of time spent waiting for CPU (latency)

Resource Controls Tab

Resource controls allow you to associate a project with a set of resource constraints. These constraints determine the allowable resource usage of tasks and processes that run in the context of the project.




FIGURE 12-2 Resource Controls Tab in the Solaris Management Console

▼ How to Access the Resource Controls Tab

The Resource Controls tab is located under System Configuration in the Navigation pane. To access Resource Controls, do the following:

1. Click the **System Configuration** control entity in the **Navigation** pane.
2. Double-click **Projects**.
3. Click on a project in the console main window to select it.
4. Select **Properties** from the **Action** menu.
5. Click the **Resource Controls** tab.

View, add, edit, or delete resource control values for processes, projects, and tasks.

Resource Controls You Can Set

To view the list of available resource controls, see About Resource Controls in the console or “Available Resource Controls” on page 90.

Setting Values

You can view, add, edit, or delete resource control values for processes, projects, and tasks. These operations are performed through dialog boxes in the console.

Resource controls and values are viewed in tables in the console. The Resource Control column lists the resource controls that can be set. The Value column displays the properties that are associated with each resource control. In the table, these values are enclosed in parentheses, and they appear as plain text separated by commas. The values in parentheses comprise an “action clause.” Each action clause is composed of a threshold, a privilege level, one signal, and one local action that is associated with the particular threshold. Each resource control can have multiple action clauses, which are also separated by commas.

Note – On a running system, values that are altered in the project database through the console only take effect for new tasks that are started in a project.

References

For information on projects and tasks, see Chapter 6. For information on resource controls, see Chapter 8. For information on the fair share scheduler (FSS), see Chapter 9.

CHAPTER **13**

Accessing Remote File Systems Topics

Chapter 14	Provides overview information for the NFS service
Chapter 15	Provides step-by-step instructions for setting up and troubleshooting the NFS service
Chapter 16	Provides background information on the NFS service

Composed September 5, 2002

Managing Network File Systems (Overview)

This chapter provides an overview of the NFS service, which can be used to access file systems over the network. The chapter includes a discussion of the concepts necessary to understand the NFS service and a description of the latest features in NFS and autofs.

- “NFS Terminology” on page 143
- “About the NFS Service” on page 144
- “About Autofs” on page 145
- “Features of the NFS Service” on page 146

NFS Terminology

This section presents some of the basic terminology that must be understood to work with the NFS service. Expanded coverage of the NFS service is included in Chapter 16.

NFS Servers and Clients

The terms *client* and *server* are used to describe the roles that a computer assumes when sharing file systems. Computers that share their file systems over a network are acting as servers. The computers that are accessing the file systems are said to be clients. The NFS service enables any computer to access any other computer’s file systems. At the same time, the NFS service provides access to its own file systems. A computer can assume the role of client, server, or both client and server at any particular time on a network.

Clients access files on the server by mounting the server’s shared file systems. When a client mounts a remote file system, the client does not make a copy of the file system. Rather, the mounting process uses a series of remote procedure calls that enable the

client to access the file system transparently on the server's disk. The mount resembles a local mount. Users type commands as if the file systems were local. See "Mounting File Systems" on page 156 for information about tasks that mount file systems.

After a file system has been shared on a server through an NFS operation, the file system can be accessed from a client. You can mount an NFS file system automatically with autofs. See "Automatic File-System Sharing" on page 152 and "Task Overview for Autofs Administration" on page 168 for tasks that involve the `share` command and autofs.

NFS File Systems

The objects that can be shared with the NFS service include any whole or partial directory tree or a file hierarchy—including a single file. A computer cannot share a file hierarchy that overlaps a file hierarchy that is already shared. Peripheral devices such as modems and printers cannot be shared.

In most UNIX system environments, a file hierarchy that can be shared corresponds to a file system or to a portion of a file system. However, NFS support works across operating systems, and the concept of a file system might be meaningless in other, non-UNIX environments. Therefore, the term *file system* refers to a file or file hierarchy that can be shared and be mounted with NFS.

About the NFS Service

The NFS service enables computers of different architectures that run different operating systems to share file systems across a network. NFS support has been implemented on many platforms that range from the MS-DOS to the VMS operating systems.

The NFS environment can be implemented on different operating systems because NFS defines an abstract model of a file system, rather than an architectural specification. Each operating system applies the NFS model to its file system semantics. This model means that file system operations such as reading and writing function as though the operations are accessing a local file.

The NFS service has the following benefits:

- Enables multiple computers to use the same files, so everyone on the network can access the same data
- Reduces storage costs by having computers share applications instead of needing local disk space for each user application

- Provides data consistency and reliability because all users can read the same set of files
- Makes mounting of file systems transparent to users
- Makes accessing of remote files transparent to users
- Supports heterogeneous environments
- Reduces system administration overhead

The NFS service makes the physical location of the file system irrelevant to the user. You can use the NFS implementation to enable users to see all the relevant files regardless of location. Instead of placing copies of commonly used files on every system, the NFS service enables you to place one copy on one computer's disk. All other systems access the files across the network. Under NFS operation, remote file systems are almost indistinguishable from local file systems.

About Autofs

File systems that are shared through the NFS service can be mounted by using automatic mounting. Autofs, a client-side service, is a file system structure that provides automatic mounting. The autofs file system is initialized by `automount`, which is run automatically when a system is booted. The `automount` daemon, `automountd`, runs continuously, mounting and unmounting remote directories as necessary.

Whenever a client computer that is running `automountd` tries to access a remote file or remote directory, the daemon mounts the remote file system. This remote file system remains mounted for as long as needed. If the remote file system is not accessed for a certain period of time, the file system is automatically unmounted.

Mounting need not be done at boot time, and the user no longer has to know the superuser password to mount a directory. Users do not need to use the `mount` and `umount` commands. The autofs service mounts and unmounts file systems as required without any intervention by the user.

Mounting some file hierarchies with `automountd` does not exclude the possibility of mounting other hierarchies with `mount`. A diskless computer *must* mount `/` (root), `/usr`, and `/usr/kvm` through the `mount` command and the `/etc/vfstab` file.

"Task Overview for Autofs Administration" on page 168 and "How Autofs Works" on page 238 give more specific information about the autofs service.

Features of the NFS Service

This section describes the important features that are included in the NFS service.

NFS Version 2 Protocol

Version 2 was the first version of the NFS protocol in wide use. Version 2 continues to be available on a large variety of platforms. All Solaris releases support version 2 of the NFS protocol, but Solaris releases prior to Solaris 2.5 support version 2 only.

NFS Version 3 Protocol

An implementation of NFS version 3 protocol was a new feature of the Solaris 2.5 release. Several changes have been made to improve interoperability and performance. For optimal use, the version 3 protocol must be running on both the NFS servers and clients.

This version enables safe asynchronous writes on the server, which improve performance by allowing the server to cache client write requests in memory. The client does not need to wait for the server to commit the changes to disk, so the response time is faster. Also, the server can batch the requests, which improves the response time on the server.

All Solaris NFS version 3 operations return the file attributes, which are stored in the local cache. Because the cache is updated more often, the need to do a separate operation to update this data arises less often. Therefore, the number of RPC calls to the server is reduced, improving performance.

The process for verifying file access permissions has been improved. Version 2 generated a “write error” message or a “read error” message if users tried to copy a remote file without the appropriate permissions. In version 3, the permissions are checked before the file is opened, so the error is reported as an “open error.”

The NFS version 3 protocol removes the 8-Kbyte transfer size limit. Clients and servers negotiate whatever transfer size the clients and servers support, rather than conform to the 8-Kbyte limit imposed in version 2. The Solaris 2.5 implementation defaults to a 32-Kbyte transfer size.

NFS ACL Support

Access control list (ACL) support was added in the Solaris 2.5 release. ACLs provide a finer-grained mechanism to set file access permissions than is available through standard UNIX file permissions. NFS ACL support provides a method of changing and viewing ACL entries from a Solaris NFS client to a Solaris NFS server. See “Using Access Control Lists (ACLs)” in *System Administration Guide: Security Services* for more information about ACLs.

NFS Over TCP

The default transport protocol for the NFS protocol was changed to the Transport Control Protocol (TCP) in the Solaris 2.5 release. TCP helps performance on slow networks and wide area networks. TCP also provides congestion control and error recovery. NFS over TCP works with version 2 and version 3. Prior to 2.5, the default NFS protocol was User Datagram Protocol (UDP).

Network Lock Manager and NFS

The Solaris 2.5 release also included an improved version of the network lock manager. The network lock manager provided UNIX record locking and PC file sharing for NFS files. The locking mechanism is now more reliable for NFS files, so commands that use locking are less likely to hang.

NFS Large File Support

The Solaris 2.6 implementation of the NFS version 3 protocol was changed to correctly manipulate files that were larger than 2 Gbytes. The NFS version 2 protocol and the Solaris 2.5 implementation of the version 3 protocol could not handle files that were larger than 2 Gbytes.

NFS Client Failover

Dynamic failover of read-only file systems was added in the Solaris 2.6 release. Failover provides a high level of availability for read-only resources that are already replicated, such as man pages, other documentation, and shared binaries. Failover can occur anytime after the file system is mounted. Manual mounts can now list multiple replicas, much like the automounter in previous releases. The automounter has not changed, except that failover need not wait until the file system is remounted. See “How to Use Client-Side Failover” on page 160 and “Client-Side Failover” on page 224 for more information.

Kerberos Support for the NFS Service

Support for Kerberos V4 clients was included in the Solaris 2.0 release. In the 2.6 release, the `mount` and `share` commands were altered to support NFS version 3 mounts that use Kerberos V5 authentication. Also, the `share` command was changed to enable multiple authentication flavors for different clients. See “RPCSEC_GSS Security Flavor” on page 148 for more information about changes that involve security flavors. See “Configuring SEAM NFS Servers” in *System Administration Guide: Security Services* for information about Kerberos V5 authentication.

WebNFS Support

The Solaris 2.6 release also included the ability to make a file system on the Internet accessible through firewalls. This capability was provided by using an extension to the NFS protocol. One of the advantages to using the WebNFS™ protocol for Internet access is its reliability. The service is built as an extension of the NFS version 3 and version 2 protocol. Also, an NFS server provides greater throughput under a heavy load than Hypertext Transfer Protocol (HTTP) access to a Web server. This throughput can decrease the amount of time that is required to retrieve a file. In addition, the WebNFS implementation provides the ability to share these files without the administrative overhead of an anonymous `ftp` site. See “Security Negotiation for the WebNFS Service” on page 149 for a description of more changes that are related to the WebNFS service. See “WebNFS Administration Tasks” on page 165 for more task information.

RPCSEC_GSS Security Flavor

A security flavor, called `RPCSEC_GSS`, is supported in the Solaris 7 release. This flavor uses the standard GSS-API interfaces to provide authentication, integrity, and privacy, as well as enabling support of multiple security mechanisms. See “Kerberos Support for the NFS Service” on page 148 for more information about support of Kerberos V5 authentication. See *GSS-API Programming Guide* for more information about GSS-API.

Solaris 7 Extensions for NFS Mounting

The Solaris 7 release includes extensions to the `mount` command and `automountd` command. The extensions enable the `mount` request to use the public file handle instead of the MOUNT protocol. The MOUNT protocol is the same access method that the WebNFS service uses. By circumventing the MOUNT protocol, the `mount` can occur through a firewall. Additionally, because fewer transactions need to occur between the server and the client, the `mount` should occur faster.

The extensions also enable NFS URLs to be used instead of the standard path name. Also, you can use the `public` option with the `mount` command and the automounter maps to force the use of the public file handle. See “WebNFS Support” on page 148 for more information about changes to the WebNFS service.

Security Negotiation for the WebNFS Service

A new protocol has been added to enable a WebNFS client to negotiate a security mechanism with an NFS server in the Solaris 8 release. This protocol provides the ability to use secure transactions when using the WebNFS service. See “How WebNFS Security Negotiation Works” on page 228 for more information.

NFS Server Logging

In the Solaris 8 release, NFS server logging enables an NFS server to provide a record of file operations that have been performed on its file systems. The record includes information on what file was accessed, when the file was accessed, and who accessed the file. You can specify the location of the logs that contain this information through a set of configuration options. You can also use these options to select the operations that should be logged. This feature is particularly useful for sites that make anonymous FTP archives available to NFS and WebNFS clients. See “How to Enable NFS Server Logging” on page 155 for more information.

Autofs Features

Autofs works with file systems that are specified in the local namespace. This information can be maintained in NIS, NIS+, or local files.

A fully multithreaded version of `automountd` was included in the Solaris 2.6 release. This enhancement makes autofs more reliable and enables concurrent servicing of multiple mounts, which prevents the service from hanging if a server is unavailable.

The new `automountd` also provides better on-demand mounting. Previous releases would mount an entire set of file systems if the file systems were hierarchically related. Now, only the top file system is mounted. Other file systems that are related to this mount point are mounted when needed.

The autofs service supports browsability of indirect maps. This support enables a user to see what directories could be mounted, without having to actually mount each one of the file systems. A `-nobrowse` option has been added to the autofs maps so that large file systems, such as `/net` and `/home`, are not automatically browsable. Also, you can turn off autofs browsability on each client by using the `-n` option with `automount`. See “Disabling Autofs Browsability” on page 180 for more information.

Composed September 5, 2002

Remote File-System Administration (Tasks)

This chapter provides information on how to perform such NFS administration tasks as setting up NFS services, adding new file systems to share, and mounting file systems. The chapter also covers the use of the Secure NFS system, and the use of WebNFS functionality. The last part of the chapter includes troubleshooting procedures and a list of some of the NFS error messages and their meanings.

- “Automatic File-System Sharing” on page 152
- “Mounting File Systems” on page 156
- “Setting Up NFS Services” on page 162
- “Administering the Secure NFS System” on page 163
- “WebNFS Administration Tasks” on page 165
- “Task Overview for Autofs Administration” on page 168
- “Strategies for NFS Troubleshooting” on page 182
- “NFS Troubleshooting Procedures” on page 183
- “NFS Error Messages” on page 191

Your responsibilities as an NFS administrator depend on your site’s requirements and the role of your computer on the network. You might be responsible for all the computers on your local network, in which instance you might be responsible for determining these configuration items:

- Which computers should be dedicated servers
- Which computers should act as both servers and clients
- Which computers should be clients only

Maintaining a server after it has been set up involves the following tasks:

- Sharing and unsharing file systems as necessary
- Modifying administrative files to update the lists of file systems your computer shares or mounts automatically
- Checking the status of the network
- Diagnosing and fixing NFS-related problems as they arise
- Setting up maps for autofs

Remember, a computer can be both a server and a client—sharing local file systems with remote computers and mounting remote file systems.

Automatic File-System Sharing

Servers provide access to their file systems by sharing the file systems over the NFS environment. You specify which file systems are to be shared with the `share` command or the `/etc/dfs/dfstab` file.

Entries in the `/etc/dfs/dfstab` file are shared automatically whenever you start NFS server operation. You should set up automatic sharing if you need to share the same set of file systems on a regular basis. For example, if your computer is a server that supports home directories, you need to make the home directories available at all times. Most file-system sharing should be done automatically. The only time that manual sharing should occur is during testing or troubleshooting.

The `dfstab` file lists all the file systems that your server shares with its clients. This file also controls which clients can mount a file system. You can modify `dfstab` to add or delete a file system or change the way sharing is done. Just edit the file with any text editor that is supported (such as `vi`). The next time that the computer enters run level 3, the system reads the updated `dfstab` to determine which file systems should be shared automatically.

Each line in the `dfstab` file consists of a `share` command—the same command that you type at the command-line prompt to share the file system. The `share` command is located in `/usr/sbin`.

TABLE 15-1 File-System Sharing Task Map

Task	Description	For Instructions
Establish automatic file-system sharing	Steps to configure a server so that file systems are automatically shared when the server is rebooted	“How to Set Up Automatic File-System Sharing” on page 152
Enable WebNFS	Steps to configure a server so that users can access files by using WebNFS	“How to Enable WebNFS Access” on page 154
Enable NFS server logging	Steps to configure a server so that NFS logging is run on selected file systems	“How to Enable NFS Server Logging” on page 155

▼ How to Set Up Automatic File-System Sharing

1. **Become superuser or assume an equivalent role.**

For information about roles, see "Using Privileged Applications" in *System Administration Guide: Security Services*.

2. Add entries for each file system to be shared.

Edit `/etc/dfs/dfstab`. Add one entry to the file for every file system that you want to be automatically shared. Each entry must be on a line by itself in the file and use this syntax:

```
share [-F nfs] [-o specific-options] [-d description] pathname
```

See the `dfstab(4)` man page for a description of `/etc/dfs/dfstab` and the `share_nfs(1M)` man page for a complete list of options.

3. Check if the NFS service is running on the server.

If this is the first `share` command or set of `share` commands that you have initiated, the NFS service might not be running. Check that one of the NFS daemons is running by using the following command.

```
# pgrep nfsd
318
```

318 is the process ID for `nfsd` in this example. If an ID is not displayed, then the service is not running. The second daemon to check for is `mountd`.

4. (Optional) Start the NFS service.

If the previous step does not report a process ID for `nfsd`, start the NFS service by using the following command.

```
# /etc/init.d/nfs.server start
```

This command ensures that NFS service is now running on the servers and restarts automatically when the server is at run level 3 during boot.

5. (Optional) Share the file system.

After the entry is in `/etc/dfs/dfstab`, the file system can be shared by either rebooting the system or by using the `shareall` command. If the NFS service was started earlier, this command does not need to be run because the init script runs the command.

```
# shareall
```

6. Verify that the information is correct.

Run the `share` command to check that the correct options are listed:

```
# share
-      /export/share/man  ro  ""
-      /usr/src           rw=eng  ""
-      /export/ftp        ro,public  ""
```

Where to Go From Here

The next step is to set up your autofs maps so that clients can access the file systems that you have shared on the server. See “Task Overview for Autofs Administration” on page 168.

▼ How to Enable WebNFS Access

Starting with the 2.6 release, by default all file systems that are available for NFS mounting are automatically available for WebNFS access. The only condition that requires the use of this procedure is one of the following:

- To allow NFS mounting on a server that does not already allow NFS mounting
- To reset the public file handle to shorten NFS URLs by using the `public` option
- To force the loading of a specific html file by using the `index` option

See “Planning for WebNFS Access” on page 166 for a list of issues that you should consider before starting the WebNFS service.

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Add entries for each file system to be shared by using the WebNFS service.

Edit `/etc/dfs/dfstab`. Add one entry to the file for every file system. The `public` and `index` tags that are shown in the following example are optional.

```
share -F nfs -o ro,public,index=index.html /export/ftp
```

See the `dfstab(4)` man page for a description of `/etc/dfs/dfstab` and the `share_nfs(1M)` man page for a complete list of options.

3. Check if the NFS service is running on the server.

If this is the first `share` command or set of `share` commands that you have initiated, the NFS daemons might not be running. Check that one of the NFS daemons is running by using the following command.

```
# pgrep nfsd
318
```

318 is the process ID for `nfsd` in this example. If an ID is not displayed, then the service is not running. The second daemon to check for is `mountd`.

4. (Optional) Start the NFS service.

If the previous step does not report a process ID for `nfsd`, start the NFS service by using the following command.

```
# /etc/init.d/nfs.server start
```

This command ensures that NFS service is now running on the servers and restarts automatically when the server is at run level 3 during boot.

5. (Optional) Share the file system.

After the entry is in `/etc/dfs/dfstab`, the file system can be shared by either rebooting the system or by using the `shareall` command. If the NFS service was started earlier, this command does not need to be run because the script runs the command.

```
# shareall
```

6. Verify that the information is correct.

Run the `share` command to check that the correct options are listed:

```
# share
-      /export/share/man    ro      ""
-      /usr/src             rw=eng  ""
-      /export/ftp          ro,public,index=index.html ""
```

▼ How to Enable NFS Server Logging

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. (Optional) Change file-system configuration settings.

In `/etc/nfs/nfslog.conf`, you can change the settings in one of two ways. You can edit the default settings for all file systems by changing the data that is associated with the `global` tag. Alternately, you can add a new tag for this file system. If these changes are not needed, you do not need to change this file. The format of `/etc/nfs/nfslog.conf` is described in `nfslog.conf(4)`.

3. Add entries for each file system to be shared by using NFS server logging.

Edit `/etc/dfs/dfstab`. Add one entry to the file for the file system on which you are enabling NFS server logging. The tag that is used with the `log=tag` option must be entered in `/etc/nfs/nfslog.conf`. This example uses the default settings in the `global` tag.

```
share -F nfs -o ro,log=global /export/ftp
```

See the `dfstab(4)` man page for a description of `/etc/dfs/dfstab` and the `share_nfs(1M)` man page for a complete list of options.

4. Check if the NFS service is running on the server.

If this is the first `share` command or set of `share` commands that you have initiated, the NFS daemons might not be running. Check that one of the NFS daemons is running by using the following command.

```
# pgrep nfsd
318
```

318 is the process ID for `nfsd` in this example. If an ID is not displayed, then the service is not running. The second daemon to check for is `mountd`.

5. (Optional) Start the NFS service.

If the previous step does not report a process ID for `nfsd`, start the NFS service by using the following command.

```
# /etc/init.d/nfs.server start
```

This command ensures that NFS service is now running on the servers and restarts automatically when the server is at run level 3 during boot.

6. (Optional) Share the file system.

After the entry is in `/etc/dfs/dfstab`, the file system can be shared by either rebooting the system or by using the `shareall` command. If the NFS service was started earlier, this command does not need to be run because the script runs the command.

```
# shareall
```

7. Verify that the information is correct.

Run the `share` command to check that the correct options are listed:

```
# share
-      /export/share/man    ro    ""
-      /usr/src             rw=eng ""
-      /export/ftp          ro,log=global ""
```

8. (Optional) Start the NFS log daemon, `nfslogd`, if it is not running already.

Restarting the NFS daemons by using the `nfs.server` script starts the daemon if the `nfslog.conf` file exists. Otherwise, the command needs to be run once by hand to create the files so that the command automatically restarts when the server is rebooted.

```
# /usr/lib/nfs/nfslogd
```

Mounting File Systems

You can mount file systems in several ways. File systems can be mounted automatically when the system is booted, on demand from the command line, or through the automounter. The automounter provides many advantages to mounting at boot time or mounting from the command line. However, many situations require a combination of all three methods. Additionally, several ways of enabling or disabling processes exist, depending on the options you use when mounting the file system. See the following table for a complete list of the tasks that are associated with file-system mounting.

TABLE 15-2 Task Map for Mounting File Systems

Task	Description	For Instructions
Mount a file system at boot time	Steps so that a file system is mounted whenever a system is rebooted.	"How to Mount a File System at Boot Time" on page 157
Mount a file system by using a command	Steps to mount a file system when a system is running. This procedure is useful when testing.	"How to Mount a File System From the Command Line" on page 158
Mount with the automounter	Steps to access a file system on demand without using the command line.	"Mounting With the Automounter" on page 159
Prevent large files	Steps to prevent large files from being created on a file system.	"How to Disable Large Files on an NFS Server" on page 159
Start client-side failover	Steps to enable the automatic switchover to a working file system if a server fails.	"How to Use Client-Side Failover" on page 160
Disable mount access for a client	Steps to disable the ability of one client to access a remote file system.	"How to Disable Mount Access for One Client" on page 160
Provide access to a file system through a firewall	Steps to allow access to a file system through a firewall by using the WebNFS protocol.	"How to Mount an NFS File System Through a Firewall" on page 161
Mount a file system by using an NFS URL	Steps to allow access to a file system by using an NFS URL. This process allows for file-system access without using the MOUNT protocol.	"How to Mount an NFS File System Using an NFS URL" on page 161

▼ How to Mount a File System at Boot Time

If you want to mount file systems at boot time instead of using autofs maps, follow this procedure. Although you must follow this procedure for all local file systems, do not use the procedure for remote file systems. This procedure must be completed on every client for remote file systems.

1. Become superuser or assume an equivalent role.

For information about roles, see "Using Privileged Applications" in *System Administration Guide: Security Services*.

2. Add an entry for the file system to `/etc/vfstab`.

Entries in the `/etc/vfstab` file have the following syntax:

```
special fsckdev mountp fstype fsckpass mount-at-boot mntopts
```

See the `vfstab(4)` man page for more information.



Caution – NFS servers should not have NFS `vfstab` entries because of a potential deadlock. The NFS service is started after the entries in `/etc/vfstab` are checked. Consider the following. If two servers that are mounting file systems from each other fail at the same time, each system could hang as the systems reboot.

Example-`vfstab` entry

You want a client computer to mount the `/var/mail` directory from the server `wasp`. You want the file system to be mounted as `/var/mail` on the client and you want the client to have read-write access. Add the following entry to the client's `vfstab` file.

```
wasp:/var/mail - /var/mail nfs - yes rw
```

▼ How to Mount a File System From the Command Line

Mounting a file system from the command line is often done to test a new mount point. This type of mount allows for temporary access to a file system that is not available through the automounter.

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Mount the file system.

Type the following command:

```
# mount -F nfs -o ro bee:/export/share/local /mnt
```

In this instance, the `/export/share/local` file system from the server `bee` is mounted on read-only `/mnt` on the local system. Mounting from the command line allows for temporary viewing of the file system. You can unmount the file system with `umount` or by rebooting the local host.



Caution – Starting with the 2.6 release, all versions of the `mount` command do not warn about invalid options. The command silently ignores any options that cannot be interpreted. To prevent unexpected behavior, ensure that you verify all of the options that were used.

Mounting With the Automounter

“Task Overview for Autofs Administration” on page 168 includes the specific instructions for establishing and supporting mounts with the automounter. Without any changes to the generic system, clients should be able to access remote file systems through the `/net` mount point. To mount the `/export/share/local` file system from the previous example, you need to type the following:

```
% cd /net/bee/export/share/local
```

Because the automounter allows all users to mount file systems, `root` access is not required. The automounter also provides for automatic unmounting of file systems, so you do not need to unmount file systems after you are finished.

▼ How to Disable Large Files on an NFS Server

For servers that are supporting clients that cannot handle a file over 2 GBytes, you might need to disable the ability to create large files.

Note – Previous versions of the Solaris operating environment cannot use large files. Check that clients of the NFS server are running at minimum the 2.6 release if the clients need to access large files.

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Check that no large files exist on the file system.

Here is an example of a command that you can run to locate large files:

```
# cd /export/home1
# find . -xdev -size +2000000 -exec ls -l {} \;
```

If large files are on the file system, you must remove or move these files to another file system.

3. Unmount the file system.

```
# umount /export/home1
```

4. Reset the file system state if the file system has been mounted by using `largefiles`.

`fsck` resets the file system state if no large files exist on the file system:

```
# fsck /export/home1
```

5. Mount the file system by using `nolargefiles`.

```
# mount -F ufs -o nolargefiles /export/home1
```

You can mount from the command line, but to make the option more permanent, add an entry that resembles the following into `/etc/vfstab`:

```
/dev/dsk/c0t3d0s1 /dev/rdsk/c0t3d0s1 /export/home1 ufs 2 yes nolargefiles
```

▼ How to Use Client-Side Failover

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. On the NFS client, mount the file system by using the `ro` option.

You can mount from the command line, through the automounter, or by adding an entry to `/etc/vfstab` that resembles the following:

```
bee,wasp:/export/share/local - /usr/local nfs - no -o ro
```

This syntax has been allowed by the automounter in earlier releases. However, the failover was not available while file systems were mounted, only when a server was being selected.

Note – Servers that are running different versions of the NFS protocol cannot be mixed by using a command line or in a `vfstab` entry. Mixing servers that support NFS V2 and V3 protocols can only be done with autofs. In autofs, the best subset of version 2 or version 3 servers is used.

▼ How to Disable Mount Access for One Client

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Add an entry in `/etc/dfs/dfstab`.

The first example allows mount access to all clients in the `eng` netgroup except the host that is named `rose`. The second example allows mount access to all clients in the `eng.example.com` DNS domain except for `rose`.

```
share -F nfs -o ro=-rose:eng /export/share/man
share -F nfs -o ro=-rose:.eng.example.com /export/share/man
```

For additional information on access lists, see “Setting Access Lists With the `share` Command” on page 213. For a description of `/etc/dfs/dfstab`, see `dfstab(4)`.

3. Share the file system.

The NFS server does not use changes to `/etc/dfs/dfstab` until the file systems are shared again or until the server is rebooted.

```
# shareall
```

▼ How to Mount an NFS File System Through a Firewall

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Manually mount the file system by using a command such as the following:

```
# mount -F nfs -o public bee:/export/share/local /mnt
```

In this example, the file system `/export/share/local` is mounted on the local client by using the `public` file handle. An NFS URL can be used instead of the standard path name. If the `public` file handle is not supported by the server `bee`, the mount operation fails.

Note – This procedure requires that the file system on the NFS server be shared by using the `public` option. Additionally, any firewalls between the client and the server must allow TCP connections on port 2049. Starting with the 2.6 release, all file systems that are shared allow for public file handle access, so the `public` option is applied by default.

▼ How to Mount an NFS File System Using an NFS URL

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Manually mount the file system by using a command such as the following:

```
# mount -F nfs nfs://bee:3000/export/share/local /mnt
```

In this example, the `/export/share/local` file system is being mounted from the server `bee` by using NFS port number 3000. The port number is not required and by default the standard NFS port number of 2049 is used. You can choose to include the `public` option with an NFS URL. Without the `public` option, the MOUNT protocol is used if the `public` file handle is not supported by the server. The `public` option forces the use of the `public` file handle, and the mount fails if the `public` file handle is

not supported.

Setting Up NFS Services

This section describes some of the tasks necessary to initialize or use NFS services.

TABLE 15–3 Task Map for NFS Services

Task	Description	For Instructions
Start the NFS server	Steps to start the NFS service if it has not been started automatically.	“How to Start the NFS Services” on page 162
Stop the NFS server	Steps to stop the NFS service. Normally the service should not need to be stopped.	“How to Stop the NFS Services” on page 162
Start the automounter	Steps to start the automounter. This procedure is required when some of the automounter maps are changed.	“How to Start the Automounter” on page 163
Stop the automounter	Steps to stop the automounter. This procedure is required when some of the automounter maps are changed.	“How to Stop the Automounter” on page 163

▼ How to Start the NFS Services

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Enable the NFS service daemons.

Type the following command:

```
# /etc/init.d/nfs.server start
```

This command starts the daemons if an entry is in `/etc/dfs/dfstab`.

▼ How to Stop the NFS Services

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Disable the NFS service daemons.

Type the following command:

```
# /etc/init.d/nfs.server stop
```

▼ How to Start the Automounter

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Enable the autofs daemon.**

Type the following command:

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

▼ How to Stop the Automounter

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Disable the autofs daemon.**

Type the following command:

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

Administering the Secure NFS System

To use the Secure NFS system, all the computers that you are responsible for must have a domain name. A domain is an administrative entity—typically, several computers—that is part of a larger network. If you are running a name service, you should also establish the name service for the domain. See *System Administration Guide: Naming and Directory Services (DNS and NIS+)*.

You can configure the Secure NFS environment to use Diffie-Hellman authentication. “Using Authentication Services (Tasks)” in *System Administration Guide: Security Services* discusses this authentication service.

Kerberos V5 authentication is also supported by the NFS service. “Introduction to SEAM” in *System Administration Guide: Security Services* discusses the Kerberos service.

▼ How to Set Up a Secure NFS Environment With DH Authentication

1. **Assign your domain a domain name, and make the domain name known to each computer in the domain.**
See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)* if you are using NIS+ as your name service.
2. **Establish public keys and secret keys for your clients' users by using the `newkey` or `nisaddcred` command. Have each user establish his or her own secure RPC password by using the `chkey` command.**

Note – For information about these commands, see the `newkey(1M)`, the `nisaddcred(1M)`, and the `chkey(1)` man pages.

When public keys and secret keys have been generated, the public keys and encrypted secret keys are stored in the `publickey` database.

3. **Verify that the name service is responding. If you are running NIS+, type the following:**

```
# nisping -u
Last updates for directory eng.acme.com. :
Master server is eng-master.acme.com.
      Last update occurred at Mon Jun  5 11:16:10 1995

Replica server is engl-replica-replica-58.acme.com.
      Last Update seen was Mon Jun  5 11:16:10 1995
```

If you are running NIS, verify that the `ypbind` daemon is running.

4. **Verify that the `keyserv` daemon of the key server is running.**

Type the following command.

```
# ps -ef | grep keyserv
root    100      1  16   Apr 11 ?        0:00 /usr/sbin/keyserv
root    2215    2211  5   09:57:28 pts/0    0:00 grep keyserv
```

If the daemon isn't running, start the key server by typing the following:

```
# /usr/sbin/keyserv
```

5. **Decrypt and store the secret key.**

Usually, the login password is identical to the network password. In this situation, `keylogin` is not required. If the passwords are different, the users have to log in, and then do a `keylogin`. You still need to use the `keylogin -r` command as root to store the decrypted secret key in `/etc/.rootkey`.

Note – You need to run `keylogin -r` if the root secret key changes or `/etc/.rootkey` is lost.

6. Update mount options for the file system.

For Diffie-Hellman authentication, edit the `/etc/dfs/dfstab` file and add the `sec=dh` option to the appropriate entries.

```
share -F nfs -o sec=dh /export/home
```

See the `dfstab(4)` man page for a description of `/etc/dfs/dfstab`.

7. Update the automounter maps for the file system.

Edit the `auto_master` data to include `sec=dh` as a mount option in the appropriate entries for Diffie-Hellman authentication:

```
/home      auto_home      -nosuid,sec=dh
```

Note – Releases through Solaris 2.5 have a limitation. If a client does not securely mount a shared file system that is secure, users have access as `nobody`, rather than as themselves. For subsequent releases that use version 2, the NFS server refuses access if the security modes do not match, unless `-sec=none` is included on the `share` command line. With version 3, the mode is inherited from the NFS server, so clients do not need to specify `sec=dh`. The users have access to the files as themselves.

When you reinstall, move, or upgrade a computer, remember to save `/etc/.rootkey` if you do not establish new keys or change the keys for `root`. If you do delete `/etc/.rootkey`, you can always type the following:

```
# keylogin -r
```

WebNFS Administration Tasks

This section provides instructions for administering the WebNFS system. Related tasks follow.

TABLE 15-4 Task Map for WebNFS Administration

Task	Description	For Instructions
Plan for WebNFS	Issues to consider before enabling the WebNFS service.	"Planning for WebNFS Access" on page 166
Enable WebNFS	Steps to enable mounting of an NFS file system by using the WebNFS protocol.	"How to Enable WebNFS Access" on page 154
Enable WebNFS through a firewall	Steps to allow access to files through a firewall by using the WebNFS protocol.	"How to Enable WebNFS Access Through a Firewall" on page 167
Browse by using an NFS URL	Instructions for using an NFS URL within a web browser.	"How to Browse Using an NFS URL" on page 167
Use a public file handle with autofs	Steps to force use of the public file handle when mounting a file system with the automounter.	"How to Use a Public File Handle With Autofs" on page 179
Use an NFS URL with autofs	Steps to add an NFS URL to the automounter maps.	"How to Use NFS URLs With Autofs" on page 180
Provide access to a file system through a firewall	Steps to allow access to a file system through a firewall by using the WebNFS protocol.	"How to Mount an NFS File System Through a Firewall" on page 161
Mount a file system by using an NFS URL	Steps to allow access to a file system by using an NFS URL. This process allows for file system access without using the MOUNT protocol.	"How to Mount an NFS File System Using an NFS URL" on page 161

Planning for WebNFS Access

To use WebNFS, you first need an application that is capable of running and loading an NFS URL (for example, `nfs://server/path`). The next step is to choose the file system that can be exported for WebNFS access. If the application is web browsing, often the document root for the web server is used. You need to consider several factors when choosing a file system to export for WebNFS access.

1. Each server has one public file handle that by default is associated with the server's root file system. The path in an NFS URL is evaluated relative to the directory with which the public file handle is associated. If the path leads to a file or directory within an exported file system, the server provides access. You can use the `public` option of the `share` command to associate the public file handle with a specific exported directory. Using this option allows URLs to be relative to the shared file system rather than to the server's root file system. The root file system does not allow web access unless the root file system is shared.
2. The WebNFS environment enables users who already have mount privileges to access files through a browser. This capability is enabled regardless of whether the file system is exported by using the `public` option. Because users already have

access to these files through the NFS setup, this access should not create any additional security risk. You only need to share a file system by using the `public` option if users who cannot mount the file system need to use WebNFS access.

3. File systems that are already open to the public make good candidates for using the `public` option. Some examples are the top directory in an ftp archive or the main URL directory for a web site.
4. You can use the `index` option with the `share` command to force the loading of an HTML file. Otherwise, you can list the directory when an NFS URL is accessed.

After a file system is chosen, review the files and set access permissions to restrict viewing of files or directories, as needed. Establish the permissions, as appropriate, for any NFS file system that is being shared. For many sites, 755 permissions for directories and 644 permissions for files provide the correct level of access.

You need to consider additional factors if both NFS and HTTP URLs are to be used to access one web site. These factors are described in “WebNFS Limitations With Web Browser Use” on page 228.

▼ How to Browse Using an NFS URL

Browsers that are capable of supporting the WebNFS service should provide access to an NFS URL that resembles the following:

```
nfs: //server[:port]/path
```

<i>server</i>	Name of the file server
<i>port</i>	Port number to use (2049, default value)
<i>path</i>	Path to file, which can be relative to the public file handle or to the root file system

Note – In most browsers, the URL service type (for example, `nfs` or `http`) is remembered from one transaction to the next. The exception occurs when a URL that includes a different service type is loaded. After you use an NFS URL, a reference to an HTTP URL might be loaded. If so, subsequent pages are loaded by using the HTTP protocol instead of the NFS protocol.

▼ How to Enable WebNFS Access Through a Firewall

You can enable WebNFS access for clients that are not part of the local subnet by configuring the firewall to allow a TCP connection on port 2049. Just allowing access for `httpd` does not allow NFS URLs to be used.

Task Overview for Autofs Administration

This section describes some of the most common tasks you might encounter in your own environment. Recommended procedures are included for each scenario to help you configure autofs to best meet your clients' needs.

Note – Use the Solstice System Management Tools or see the *System Administration Guide: Naming and Directory Services (FNS and NIS+)* to perform the tasks that are discussed in this section.

Task Map for Autofs Administration

The following table provides a description and a pointer to many of the tasks that are related to autofs.

TABLE 15–5 Task Map for Autofs Administration

Task	Description	For Instructions
Start autofs	Start the automount service without having to reboot the system	"How to Start the Automounter" on page 163
Stop autofs	Stop the automount service without disabling other network services	"How to Stop the Automounter" on page 163
Access file systems by using autofs	Access file systems by using the automount service	"Mounting With the Automounter" on page 159
Modify the autofs maps	Steps to modify the master map, which should be used to list other maps	"How to Modify the Master Map" on page 171
	Steps to modify an indirect map, which should be used for most maps	"How to Modify Indirect Maps" on page 171
	Steps to modify a direct map, which should be used when a direct association between a mount point on a client and a server is required	"How to Modify Direct Maps" on page 171
Modify the autofs maps to access non-NFS file systems	Steps to set up an autofs map with an entry for a CD-ROM application	"How to Access CD-ROM Applications With Autofs" on page 173

TABLE 15-5 Task Map for Autofs Administration (Continued)

Task	Description	For Instructions
	Steps to set up an autofs map with an entry for a PC-DOS diskette	"How to Access PC-DOS Data Diskettes With Autofs" on page 173
	Steps to use autofs to access a CacheFS file system	"How to Access NFS File Systems Using CacheFS" on page 174
Using /home	Example of how to set up a common /home map	"Setting Up a Common View of /home" on page 174
	Steps to set up a /home map that refers to multiple file systems	"How to Set Up /home With Multiple Home Directory File Systems" on page 175
Using a new autofs mount point	Steps to set up a project-related autofs map	"How to Consolidate Project-Related Files Under /ws" on page 176
	Steps to set up an autofs map that supports different client architectures	"How to Set Up Different Architectures to Access a Shared Namespace" on page 177
	Steps to set up an autofs map that supports different operating systems	"How to Support Incompatible Client Operating System Versions" on page 178
Replicate file systems with autofs	Provide access to file systems that fail over	"How to Replicate Shared Files Across Several Servers" on page 179
Using security restrictions with autofs	Provide access to file systems while restricting remote root access to the files	"How to Apply Autofs Security Restrictions" on page 179
Using a public file handle with autofs	Force use of the public file handle when mounting a file system	"How to Use a Public File Handle With Autofs" on page 179
Using an NFS URL with autofs	Add an NFS URL so that the automounter can use it	"How to Use NFS URLs With Autofs" on page 180
Disable autofs browsability	Steps to disable browsability so that autofs mount points are not automatically populated on a single client	"How to Completely Disable Autofs Browsability on a Single NFS Client" on page 180
	Steps to disable browsability so that autofs mount points are not automatically populated on all clients	"How to Disable Autofs Browsability for All Clients" on page 181
	Steps to disable browsability so that a specific autofs mount point is not automatically populated on a client	"How to Disable Autofs Browsability on a Selected File System" on page 181

Administrative Tasks Involving Maps

The following tables describe several of the factors you need to be aware of when administering autofs maps. Which type of map and which name service you choose change the mechanism that you need to use to make changes to the autofs maps.

The following table describes the types of maps and their uses.

TABLE 15-6 Types of autofs Maps and Their Uses

Type of Map	Use
Master	Associates a directory with a map
Direct	Directs autofs to specific file systems
Indirect	Directs autofs to reference-oriented file systems

The following table describes how to make changes to your autofs environment that are based on your name service.

TABLE 15-7 Map Maintenance

Name Service	Method
Local files	Text editor
NIS	make files
NIS+	nistbladm

The next table tells you when to run the `automount` command, depending on the modification you have made to the type of map. For example, if you have made an addition or a deletion to a direct map, you need to run the `automount` command on the local system. By running the command, you make the change effective. However, if you have modified an existing entry, you do not need to run the `automount` command for the change to become effective.

TABLE 15-8 When to Run the `automount` Command

Type of Map	Restart <code>automount</code> ?	
	Addition or Deletion	Modification
<code>auto_master</code>	Y	Y
<code>direct</code>	Y	N
<code>indirect</code>	N	N

Modifying the Maps

The following procedures require that you use NIS+ as your name service.

▼ How to Modify the Master Map

1. **Log in as a user who has permissions to change the maps.**
2. **Using the `nistbladm` command, make your changes to the master map.**
See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)*.
3. **For each client, become superuser or assume an equivalent role.**
For information about roles, see "Using Privileged Applications" in *System Administration Guide: Security Services*.
4. **For each client, run the `automount` command to ensure that your changes become effective.**
5. **Notify your users of the changes.**
Notification is required so that the users can also run the `automount` command as superuser on their own computers.

The `automount` command gathers information from the master map whenever it is run.

▼ How to Modify Indirect Maps

1. **Log in as a user who has permissions to change the maps.**
2. **Using the `nistbladm` command, make your changes to the indirect map.**
See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)*.

The change becomes effective the next time that the map is used, which is the next time a mount is done.

▼ How to Modify Direct Maps

1. **Log in as a user who has permissions to change the maps.**
2. **Using the `nistbladm` command, add or delete your changes to the direct map.**
See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)*.
3. **If you added or deleted a mount-point entry in step 1, run the `automount` command.**

4. Notify your users of the changes.

Notification is required so that the users can also run the automount command as superuser on their own computers.

Note – If you only modify or change the contents of an existing direct map entry, you do not need to run the automount command.

For example, suppose you modify the `auto_direct` map so that the `/usr/src` directory is now mounted from a different server. If `/usr/src` is not mounted at this time, the new entry becomes effective immediately when you try to access `/usr/src`. If `/usr/src` is mounted now, you can wait until the auto-unmounting occurs, then access the file.

Note – Use indirect maps whenever possible. Indirect maps are easier to construct and less demanding on the computers' file systems. Also, indirect maps do not occupy as much space in the mount table as direct maps.

Avoiding Mount-Point Conflicts

If you have a local disk partition that is mounted on `/src` and you plan to use the autofs service to mount other source directories, you might encounter a problem. If you specify the mount point `/src`, the NFS service hides the local partition whenever you try to reach it.

You need to mount the partition in some other location, for example, on `/export/src`. You then need an entry in `/etc/vfstab` such as the following:

```
/dev/dsk/d0t3d0s5 /dev/rdsk/c0t3d0s5 /export/src ufs 3 yes -
```

You also need this entry in `auto_src`:

```
terra          terra:/export/src
```

terra is the name of the computer.

Accessing Non-NFS File Systems

Autofs can also mount files other than NFS files. Autofs mounts files on removable media, such as diskettes or CD-ROM. Normally, you would mount files on removable media by using the Volume Manager. The following examples show how this mounting could be accomplished through autofs. The Volume Manager and autofs do not work together, so these entries would not be used without first deactivating the Volume Manager.

Instead of mounting a file system from a server, you put the media in the drive and reference the file system from the map. If you plan to access non-NFS file systems and you are using autofs, see the following procedures.

How to Access CD-ROM Applications With Autofs

Note – Use this procedure if you are *not* using Volume Manager.

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Update the autofs map.

Add an entry for the CD-ROM file system, which should resemble the following:

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

The CD-ROM device that you intend to mount must appear as a name that follows the colon.

▼ How to Access PC-DOS Data Diskettes With Autofs

Note – Use this procedure if you are *not* using Volume Manager.

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Update the autofs map.

Add an entry for the diskette file system such as the following:

```
pcfs      -fstype=pcfs      :/dev/diskette
```

Accessing NFS File Systems Using CacheFS

The cache file system (CacheFS) is a generic nonvolatile caching mechanism. CacheFS improves the performance of certain file systems by utilizing a small, fast local disk.

You can improve the performance of the NFS environment by using CacheFS to cache data from an NFS file system on a local disk.

▼ How to Access NFS File Systems Using CacheFS

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Run the `cfsadmin` command to create a cache directory on the local disk.**

```
# cfsadmin -c /var/cache
```

3. **Add the `cacheafs` entry to the appropriate automounter map.**

For example, adding this entry to the master map caches all home directories:

```
/home auto_home -fstype=cacheafs,cachedir=/var/cache,backfstype=nfs
```

Adding this entry to the `auto_home` map only caches the home directory for the user who is named `rich`:

```
rich -fstype=cacheafs,cachedir=/var/cache,backfstype=nfs dragon:/export/home1/rich
```

Note – Options that are included in maps that are searched later override options set in maps that are searched earlier. The last options that are found are the ones that are used. In the previous example, an additional entry to the `auto_home` map only needs to include the options in the master maps if some options required changes.

Customizing the Automounter

You can set up the automounter maps in several ways. The following tasks give details on how to customize the automounter maps to provide an easy-to-use directory structure.

▼ Setting Up a Common View of `/home`

The ideal is for all network users to be able to locate their own or anyone’s home directory under `/home`. This view should be common across all computers, whether client or server.

Every Solaris installation comes with a master map: `/etc/auto_master`.

```
# Master map for autofs
#
+auto_master
/net      -hosts      -nosuid,nobrowse
/home     auto_home   -nobrowse
/xfn      -xfn
```

A map for `auto_home` is also installed under `/etc`.

```
# Home directory map for autofs
#
+auto_home
```

Except for a reference to an external `auto_home` map, this map is empty. If the directories under `/home` are to be common to all computers, do not modify this `/etc/auto_home` map. All home directory entries should appear in the name service files, either NIS or NIS+.

Note – Users should not be permitted to run `setuid` executables from their home directories. Without this restriction, any user could have superuser privileges on any computer.

▼ How to Set Up `/home` With Multiple Home Directory File Systems

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Install home directory partitions under `/export/home`.

If the system has several partitions, install the partitions under separate directories, for example, `/export/home1`, `/export/home2`, and so on.

3. Use the Solstice System Management Tools to create and maintain the `auto_home` map.

Whenever you create a new user account, type the location of the user’s home directory in the `auto_home` map. Map entries can be simple, for example:

```
rusty      dragon:/export/home1/&
gwenda     dragon:/export/home1/&
charles    sundog:/export/home2/&
rich       dragon:/export/home3/&
```

Notice the use of the `&` (ampersand) to substitute the map key. The ampersand is an abbreviation for the second occurrence of `rusty` in the following example.

```
rusty      dragon:/export/home1/rusty
```

With the `auto_home` map in place, users can refer to any home directory (including their own) with the path `/home/user`. *user* is their login name and the key in the map. This common view of all home directories is valuable when logging in to another user’s computer. Autofs mounts your home directory for you. Similarly, if you run a remote windowing system client on another computer, the client program has the same view of the `/home` directory.

This common view also extends to the server. Using the previous example, if `rusty` logs in to the server `dragon`, autofs there provides direct access to the local disk by

loopback-mounting /export/home1/rusty onto /home/rusty.

Users do not need to be aware of the real location of their home directories. If rusty needs more disk space and needs to have his home directory relocated to another server, you need only change rusty's entry in the auto_home map to reflect the new location. Other users can continue to use the /home/rusty path.

▼ How to Consolidate Project-Related Files Under /ws

Assume you are the administrator of a large software development project. You plan to make all project-related files available under a directory that is called /ws. This directory is to be common across all workstations at the site.

1. **Add an entry for the /ws directory to the site auto_master map, either NIS or NIS+.**

```
/ws      auto_ws      -nosuid
```

The auto_ws map determines the contents of the /ws directory.

2. **Add the -nosuid option as a precaution.**

This option prevents users from running setuid programs that might exist in any workspaces.

3. **Add entries to the auto_ws map.**

The auto_ws map is organized so that each entry describes a subproject. Your first attempt yields a map that resembles the following:

```
compiler  alpha:/export/ws/&
windows   alpha:/export/ws/&
files     bravo:/export/ws/&
drivers   alpha:/export/ws/&
man       bravo:/export/ws/&
tools     delta:/export/ws/&
```

The ampersand (&) at the end of each entry is an abbreviation for the entry key. For instance, the first entry is equivalent to the following:

```
compiler      alpha:/export/ws/compiler
```

This first attempt provides a map that appears simple, but the map is inadequate. The project organizer decides that the documentation in the man entry should be provided as a subdirectory under each subproject. Also, each subproject requires subdirectories to describe several versions of the software. You must assign each of these subdirectories to an entire disk partition on the server.

Modify the entries in the map as follows:

```
compiler \
/vers1.0   alpha:/export/ws/&/vers1.0 \
/vers2.0   bravo:/export/ws/&/vers2.0 \
```



```

    /man      bravo:/export/ws/&/man
windows \
  /vers1.0   alpha:/export/ws/&/vers1.0 \
  /man       bravo:/export/ws/&/man
files \
  /vers1.0   alpha:/export/ws/&/vers1.0 \
  /vers2.0   bravo:/export/ws/&/vers2.0 \
  /vers3.0   bravo:/export/ws/&/vers3.0 \
  /man       bravo:/export/ws/&/man
drivers \
  /vers1.0   alpha:/export/ws/&/vers1.0 \
  /man       bravo:/export/ws/&/man
tools \
  /          delta:/export/ws/&

```

Although the map now appears to be much larger, the map still contains only the five entries. Each entry is larger because each entry contains multiple mounts. For instance, a reference to `/ws/compiler` requires three mounts for the `vers1.0`, `vers2.0`, and `man` directories. The backslash at the end of each line informs autofs that the entry is continued onto the next line. Effectively, the entry is one long line, though line breaks and some indenting have been used to make the entry more readable. The `tools` directory contains software development tools for all subprojects, so this directory is not subject to the same subdirectory structure. The `tools` directory continues to be a single mount.

This arrangement provides the administrator with much flexibility. Software projects typically consume substantial amounts of disk space. Through the life of the project, you might be required to relocate and expand various disk partitions. If these changes are reflected in the `auto_ws` map, the users do not need to be notified, as the directory hierarchy under `/ws` is not changed.

Because the servers `alpha` and `bravo` view the same autofs map, any users who log in to these computers can find the `/ws` name space as expected. These users are provided with direct access to local files through loopback mounts instead of NFS mounts.

▼ How to Set Up Different Architectures to Access a Shared Namespace

You need to assemble a shared name space for local executables, and applications, such as spreadsheet applications and word-processing packages. The clients of this namespace use several different workstation architectures that require different executable formats. Also, some workstations are running different releases of the operating system.

1. **Create the `auto_local` map with the `nistbladm` command.**
See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)*.
2. **Choose a single, site-specific name for the shared namespace. This name makes the files and directories that belong to this space easily identifiable.**

For example, if you choose `/usr/local` as the name, the path `/usr/local/bin` is obviously a part of this name space.

3. **For ease of user community recognition, create an autofs indirect map. Mount this map at `/usr/local`. Set up the following entry in the NIS+ (or NIS) `auto_master` map:**

```
/usr/local      auto_local      -ro
```

Notice that the `-ro` mount option implies that clients cannot write to any files or directories.

4. **Export the appropriate directory on the server.**

5. **Include a `bin` entry in the `auto_local` map.**

Your directory structure resembles the following:

```
bin      aa:/export/local/bin
```

6. **(Optional) To serve clients of different architectures, change the entry by adding the autofs CPU variable.**

```
bin      aa:/export/local/bin/$CPU
```

- For SPARC clients – Place executables in `/export/local/bin/sparc`.
- For IA clients – Place executables in `/export/local/bin/i386`.

▼ How to Support Incompatible Client Operating System Versions

1. **Combine the architecture type with a variable that determines the operating system type of the client.**

You can combine the autofs `OSREL` variable with the `CPU` variable to form a name that determines both CPU type and OS release.

2. **Create the following map entry.**

```
bin      aa:/export/local/bin/$CPU$OSREL
```

For clients that are running version 5.6 of the operating system, export the following file systems:

- For SPARC clients – Export `/export/local/bin/sparc5.6`.
- For IA clients – Place executables in `/export/local/bin/i3865.6`.

▼ How to Replicate Shared Files Across Several Servers

The best way to share replicated file systems that are read-only is to use failover. See “Client-Side Failover” on page 224 for a discussion of failover.

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Modify the entry in the autofs maps.**

Create the list of all replica servers as a comma-separated list, such as the following:

```
bin      aa,bb,cc,dd:/export/local/bin/$CPU
```

Autofs chooses the nearest server. If a server has several network interfaces, list each interface. Autofs chooses the nearest interface to the client, avoiding unnecessary routing of NFS traffic.

▼ How to Apply Autofs Security Restrictions

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Create the following entry in the name service `auto_master` file, either NIS or NIS+:**

```
/home      auto_home      -nosuid
```

The `nosuid` option prevents users from creating files with the `setuid` or `setgid` bit set.

This entry overrides the entry for `/home` in a generic local `/etc/auto_master` file. See the previous example. The override happens because the `+auto_master` reference to the external name service map occurs before the `/home` entry in the file. If the entries in the `auto_home` map include mount options, the `nosuid` option is overwritten. Therefore, either no options should be used in the `auto_home` map or the `nosuid` option must be included with each entry.

Note – Do not mount the home directory disk partitions on or under `/home` on the server.

▼ How to Use a Public File Handle With Autofs

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Create an entry in the autofs map such as the following:

```
/usr/local    -ro,public    bee:/export/share/local
```

The `public` option forces the public handle to be used. If the NFS server does not support a public file handle, the mount fails.

▼ How to Use NFS URLs With Autofs

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Create an autofs entry such as the following:

```
/usr/local    -ro    nfs://bee/export/share/local
```

The service tries to use the public file handle on the NFS server. However, if the server does not support a public file handle, the MOUNT protocol is used.

Disabling Autofs Browsability

Starting with the Solaris 2.6 release, the default version of `/etc/auto_master` that is installed has the `-nobrowse` option added to the entries for `/home` and `/net`. In addition, the upgrade procedure adds the `-nobrowse` option to the `/home` and `/net` entries in `/etc/auto_master` if these entries have not been modified. However, you might have to make these changes manually or to turn off browsability for site-specific autofs mount points after the installation.

You can turn off the browsability feature in several ways. Disable the feature by using a command-line option to the `automountd` daemon, which completely disables autofs browsability for the client. Or disable browsability for each map entry on all clients by using the autofs maps in either an NIS or NIS+ name space. You can also disable the feature for each map entry on each client, using local autofs maps if no network-wide namespace is being used.

▼ How to Completely Disable Autofs Browsability on a Single NFS Client

1. Become superuser or assume an equivalent role on the NFS client.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Add the `-n` option to the startup script.

As root, edit the `/etc/init.d/autofs` script and add the `-n` option to the line that starts the `autmountd` daemon:

```
/usr/lib/autofs/automountd -n \  
    < /dev/null > /dev/console 2>&1 # start daemon
```

3. Restart the autofs service.

```
# /etc/init.d/autofs stop  
# /etc/init.d/autofs start
```

▼ How to Disable Autofs Browsability for All Clients

To disable browsability for all clients, you must employ a name service such as NIS or NIS+. Otherwise, you need to manually edit the automounter maps on each client. In this example, the browsability of the `/home` directory is disabled. You must follow this procedure for each indirect autofs node that needs to be disabled.

1. Add the `-nobrowse` option to the `/home` entry in the name service `auto_master` file.

```
/home      auto_home      -nobrowse
```

2. Run the `automount` command on all clients.

The new behavior becomes effective after you run the `automount` command on the client systems or after a reboot.

```
# /usr/sbin/automount
```

▼ How to Disable Autofs Browsability on a Selected File System

In this example, browsability of the `/net` directory is disabled. You can use the same procedure for `/home` or any other autofs mount points.

1. Check the `automount` entry in `/etc/nsswitch.conf`.

For local file entries to have precedence, the entry in the name service switch file should list files before the name service. For example:

```
automount:  files nisplus
```

This entry shows the default configuration in a standard Solaris installation.

2. Check the position of the `+auto_master` entry in `/etc/auto_master`.

For additions to the local files to have precedence over the entries in the namespace, the `+auto_master` entry must be moved to follow `/net`:

```
# Master map for automounter
#
/net      -hosts      -nosuid
/home     auto_home
/xfn      -xfn
+auto_master
```

A standard configuration places the `+auto_master` entry at the top of the file. This placement prevents any local changes from being used.

3. Add the `nobrowse` option to the `/net` entry in the `/etc/auto_master` file.

```
/net      -hosts      -nosuid,nobrowse
```

4. On all clients, run the `automount` command.

The new behavior becomes effective after running the `automount` command on the client systems or after a reboot.

```
# /usr/sbin/automount
```

Strategies for NFS Troubleshooting

When tracking down an NFS problem, remember the main points of possible failure: the server, the client, and the network. The strategy that is outlined in this section tries to isolate each individual component to find the one that is not working. In all situations, the `mountd` and `nfstd` daemons must be running on the server in order for remote mounts to succeed.

Note – The `mountd` and `nfstd` daemons start automatically at boot time only if NFS share entries are in the `/etc/dfs/dfstab` file. Therefore, you must start `mountd` and `nfstd` manually when you set up sharing for the first time.

The `-intr` option is set by default for all mounts. If a program hangs with a “server not responding” message, you can kill the program with the keyboard interrupt `Control-c`.

When the network or server has problems, programs that access hard-mounted remote files fail differently than those programs that access soft-mounted remote files. Hard-mounted remote file systems cause the client’s kernel to retry the requests until the server responds again. Soft-mounted remote file systems cause the client’s system calls to return an error after trying for awhile. Because these errors can result in unexpected application errors and data corruption, avoid soft mounting.

When a file system is hard mounted, a program that tries to access the file system hangs if the server fails to respond. In this situation, the NFS system displays the following message on the console:

```
NFS server hostname not responding still trying
```

When the server finally responds, the following message appears on the console:

```
NFS server hostname ok
```

A program that accesses a soft-mounted file system whose server is not responding generates the following message:

```
NFS operation failed for server hostname: error # (error_message)
```

Note – Because of possible errors, do not soft-mount file systems with read-write data or file systems from which executables are run. Writable data could be corrupted if the application ignores the errors. Mounted executables might not load properly and can fail.

NFS Troubleshooting Procedures

To determine where the NFS service has failed, you need to follow several procedures to isolate the failure. Check for the following items:

- Can the client reach the server?
- Can the client contact the NFS services on the server?
- Are the NFS services running on the server?

In the process of checking these items, you might notice that other portions of the network are not functioning. For example, the name service or the physical network hardware might not be functioning. The *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)* contains debugging procedures for several name services. Also, during the process you might see that the problem is not at the client end. An example is if you get at least one trouble call from every subnet in your work area. In this situation, you should assume that the problem is the server or the network hardware near the server. So, you should start the debugging process at the server, not at the client.

▼ How to Check Connectivity on an NFS Client

1. Check that the NFS server is reachable from the client. On the client, type the following command.

```
% /usr/sbin/ping bee
bee is alive
```

If the command reports that the server is alive, remotely check the NFS server. See “How to Check the NFS Server Remotely” on page 184.

2. If the server is not reachable from the client, ensure that the local name service is running.

For NIS+ clients, type the following:

```
% /usr/lib/nis/nisping -u
Last updates for directory eng.acme.com. :
Master server is eng-master.acme.com.
      Last update occurred at Mon Jun  5 11:16:10 1995

Replica server is eng1-replica-58.acme.com.
      Last Update seen was Mon Jun  5 11:16:10 1995
```

3. If the name service is running, ensure that the client has received the correct host information by typing the following:

```
% /usr/bin/getent hosts bee
129.144.83.117      bee.eng.acme.com
```

4. If the host information is correct, but the server is not reachable from the client, run the ping command from another client.

If the command run from a second client fails, see “How to Verify the NFS Service on the Server” on page 185.

5. If the server is reachable from the second client, use ping to check connectivity of the first client to other systems on the local net.

If this command fails, check the networking software configuration on the client (/etc/netmasks, /etc/nsswitch.conf, and so forth).

6. If the software is correct, check the networking hardware.

Try moving the client onto a second net drop.

▼ How to Check the NFS Server Remotely

1. Check that the NFS services have started on the NFS server by typing the following command:

```
% rpcinfo -s bee | egrep 'nfs|mountd'
100003  3,2      tcp,udp,tcp6,udp6      nfs      superuser
100005  3,2,1      ticots,ticotsord,tcp,tcp6,ticlts,udp,udp6 mountd  superuser
```

If the daemons have not been started, see “How to Restart NFS Services” on page 187.

2. Check that the server's nfsd processes are responding.

On the client, type the following command to test the UDP NFS connections from the server.


```
% /usr/bin/rpcinfo -u bee nfs
program 100003 version 2 ready and waiting
program 100003 version 3 ready and waiting
```

If the server is running, it prints a list of program and version numbers. Using the `-t` option tests the TCP connection. If this command fails, proceed to “How to Verify the NFS Service on the Server” on page 185.

3. Check that the server’s `mountd` is responding, by typing the following command.

```
% /usr/bin/rpcinfo -u bee mountd
program 100005 version 1 ready and waiting
program 100005 version 2 ready and waiting
program 100005 version 3 ready and waiting
```

If the server is running, it prints a list of program and version numbers that are associated with the UDP protocol. Using the `-t` option tests the TCP connection. If either attempt fails, proceed to “How to Verify the NFS Service on the Server” on page 185.

4. Check the local `autofs` service if it is being used:

```
% cd /net/wasp
```

Choose a `/net` or `/home` mount point that you know should work properly. If this command fails, then as `root` on the client, type the following to restart the `autofs` service:

```
# /etc/init.d/autofs stop
# /etc/init.d/autofs start
```

5. Verify that file system is shared as expected on the server.

```
% /usr/sbin/showmount -e bee
/usr/src                               eng
/export/share/man                      (everyone)
```

Check the entry on the server and the local mount entry for errors. Also, check the namespace. In this instance, if the first client is not in the `eng` netgroup, that client cannot mount the `/usr/src` file system.

Check all entries that include mounting information in all of the local files. The list includes `/etc/vfstab` and all the `/etc/auto_*` files.

▼ How to Verify the NFS Service on the Server

1. Become superuser or assume an equivalent role.

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Check that the server can reach the clients.

```
# ping lilac
lilac is alive
```

3. If the client is not reachable from the server, ensure that the local name service is running. For NIS+ clients, type the following:

```
% /usr/lib/nis/nisping -u
Last updates for directory eng.acme.com. :
Master server is eng-master.acme.com.
    Last update occurred at Mon Jun  5 11:16:10 1995

Replica server is eng1-replica-58.acme.com.
    Last Update seen was Mon Jun  5 11:16:10 1995
```

4. If the name service is running, check the networking software configuration on the server (/etc/netmasks, /etc/nsswitch.conf, and so forth).

5. Type the following command to check whether the `nfsd` daemon is running.

```
# rpcinfo -u localhost nfs
program 100003 version 2 ready and waiting
program 100003 version 3 ready and waiting
# ps -ef | grep nfsd
root    232      1  0 Apr 07    ?        0:01 /usr/lib/nfs/nfsd -a 16
root    3127    2462  1 09:32:57 pts/3    0:00 grep nfsd
```

If the server is running, it prints a list of program and version numbers that are associated with the UDP protocol. Also use the `-t` option with `rpcinfo` to check the TCP connection. If these commands fail, restart the NFS service. See “How to Restart NFS Services” on page 187.

6. Type the following command to check whether the `mountd` daemon is running.

```
# /usr/bin/rpcinfo -u localhost mountd
program 100005 version 1 ready and waiting
program 100005 version 2 ready and waiting
program 100005 version 3 ready and waiting
# ps -ef | grep mountd
root    145      1  0 Apr 07    ?        21:57 /usr/lib/autofs/automountd
root    234      1  0 Apr 07    ?        0:04 /usr/lib/nfs/mountd
root    3084    2462  1 09:30:20 pts/3    0:00 grep mountd
```

If the server is running, it prints a list of program and version numbers that are associated with the UDP protocol. Also use the `-t` option with `rpcinfo` to check the TCP connection. If these commands fail, restart the NFS service. See “How to Restart NFS Services” on page 187.

7. Type the following command to check whether the `rpcbind` daemon is running.

```
# /usr/bin/rpcinfo -u localhost rpcbind
program 100000 version 1 ready and waiting
program 100000 version 2 ready and waiting
program 100000 version 3 ready and waiting
```

If the server is running, it prints a list of program and version numbers that are associated with the UDP protocol. If `rpcbind` seems to be hung, either reboot the server or follow the steps in “How to Warm-Start `rpcbind`” on page 187.

▼ How to Restart NFS Services

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **To enable daemons without rebooting, type the following commands.**

```
# /etc/init.d/nfs.server stop
# /etc/init.d/nfs.server start
```

This remedy stops and restarts the daemons if an entry is in `/etc/dfs/dfstab`.

▼ How to Warm-Start rpcbind

If the NFS server cannot be rebooted because of work in progress, you can restart `rpcbind` without having to restart all of the services that use RPC. Just complete a warm start as described in this procedure.

1. **Become superuser or assume an equivalent role.**

For information about roles, see “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Determine the PID for `rpcbind`.**

Run `ps` to get the PID, which is the value in the second column.

```
# ps -ef |grep rpcbind
root    115      1  0   May 31 ?          0:14 /usr/sbin/rpcbind
root 13000  6944  0 11:11:15 pts/3    0:00 grep rpcbind
```

3. **Send a `SIGTERM` signal to the `rpcbind` process.**

In this example, `term` is the signal that is to be sent and 115 is the PID for the program (see the `kill(1)` man page). This command causes `rpcbind` to create a list of the current registered services in `/tmp/portmap.file` and `/tmp/rpcbind.file`.

```
# kill -s term 115
```

Note – If you do not kill the `rpcbind` process with the `-s term` option, you cannot complete a warm start of `rpcbind`. You must reboot the server to restore service.

4. **Restart `rpcbind`.**

Warm-restart the command so that the files that were created by the `kill` command are consulted. A warm-start also ensures that the process resumes without requiring a restart of all of the RPC services. See the `rpcbind(1M)` man page.

```
# /usr/sbin/rpcbind -w
```

▼ Identifying Which Host Is Providing NFS File Service

Run the `nfsstat` command with the `-m` option to gather current NFS information. The name of the current server is printed after “`currserver=`”.

```
% nfsstat -m
/usr/local from bee, wasp:/export/share/local
Flags: vers=3,proto=tcp,sec=sys,hard,intr,llock,link,synlink,
      acl,rsz=32768,wsz=32678,retrans=5
Failover: noresponse=0, failover=0, remap=0, currserver=bee
```

▼ How to Verify Options Used With the mount Command

In the Solaris 2.6 release and in any versions of the `mount` command that were patched after the 2.6 release, no warning is issued for invalid options. The following procedure helps determine whether the options that were supplied either on the command line or through `/etc/vfstab` were valid.

For this example, assume that the following command has been run:

```
# mount -F nfs -o ro,vers=2 bee:/export/share/local /mnt
```

1. Verify the options by running the following command.

```
% nfsstat -m
/mnt from bee:/export/share/local
Flags: vers=2,proto=tcp,sec=sys,hard,intr,dynamic,acl,rsz=8192,wsz=8192,
      retrans=5
```

The file system from `bee` has been mounted with the protocol version set to 2. Unfortunately, the `nfsstat` command does not display information about all of the options. However, using the `nfsstat` command is the most accurate way to verify the options.

2. Check the entry in `/etc/mnttab`.

The `mount` command does not allow invalid options to be added to the mount table. Therefore, verify that the options that are listed in the file match those options that are listed on the command line. In this way, you can check those options that are not reported by the `nfsstat` command.

```
# grep bee /etc/mnttab
bee:/export/share/local /mnt nfs    ro,vers=2,dev=2b0005e 859934818
```

Troubleshooting Autofs

Occasionally, you might encounter problems with autofs. This section should improve the problem-solving process. The section is divided into two subsections.

This section presents a list of the error messages that autofs generates. The list is divided into two parts:

- Error messages that are generated by the verbose (-v) option of automount
- Error messages that might appear at any time

Each error message is followed by a description and probable cause of the message.

When troubleshooting, start the autofs programs with the verbose (-v) option. Otherwise, you might experience problems without knowing why.

The following paragraphs are labeled with the error message you are likely to see if autofs fails, and a description of the possible problem.

Error Messages Generated by automount -v

`bad key key in direct map mapname`

While scanning a direct map, autofs has found an entry key without a prefixed /. Keys in direct maps must be full path names.

`bad key key in indirect map mapname`

While scanning an indirect map, autofs has found an entry key that contains a /. Indirect map keys must be simple names—not path names.

`can't mount server:pathname: reason`

The mount daemon on the server refuses to provide a file handle for *server:pathname*. Check the export table on the server.

`couldn't create mount point mountpoint: reason`

Autofs was unable to create a mount point that was required for a mount. This problem most frequently occurs when you attempt to hierarchically mount all of a server's exported file systems. A required mount point can exist only in a file system that cannot be mounted, which means the file system cannot be exported. The mount point cannot be created because the exported parent file system is exported read-only.

`leading space in map entry entry text in mapname`

Autofs has discovered an entry in an automount map that contains leading spaces. This problem is usually an indication of an improperly continued map entry. For example:

```
fake
/blat          frobz:/usr/frotz
```

In this example, the warning is generated when autofs encounters the second line because the first line should be terminated with a backslash (\).

mapname: Not found

The required map cannot be located. This message is produced only when the *-v* option is used. Check the spelling and path name of the map name.

remount *server:pathname* on *mountpoint*: server not responding

Autofs has failed to remount a file system that it previously unmounted.

WARNING: *mountpoint* already mounted on

Autofs is attempting to mount over an existing mount point. This message means an internal error occurred in autofs (an anomaly).

Miscellaneous Error Messages

dir *mountpoint* must start with '/'

The automounter mount point must be given as a full path name. Check the spelling and path name of the mount point.

hierarchical mountpoints: *pathname1* and *pathname2*

Autofs does not allow its mount points to have a hierarchical relationship. An autofs mount point must not be contained within another automounted file system.

host *server* not responding

Autofs attempted to contact *server*, but received no response.

hostname: exports: *rpc_err*

An error occurred while getting the export list from *hostname*. This message indicates a server or network problem.

map *mapname*, key *key*: bad

The map entry is malformed, and autofs cannot interpret the entry. Recheck the entry. Perhaps the entry has characters that need escaping.

mapname: *nis_err*

An error occurred when looking up an entry in a NIS map. This message can indicate NIS problems.

mount of *server:pathname* on *mountpoint*:*reason*

Autofs failed to do a mount. This occurrence can indicate a server or network problem.

mountpoint: Not a directory

Autofs cannot mount itself on *mountpoint* because it is not a directory. Check the spelling and path name of the mount point.

nfscast: cannot send packet: *reason*

Autofs cannot send a query packet to a server in a list of replicated file system locations.

nfscast: cannot receive reply: *reason*

Autofs cannot receive replies from any of the servers in a list of replicated file system locations.

nfscast: select: *reason*

All these error messages indicate problems in attempting to ping servers for a replicated file system. This message can indicate a network problem.

pathconf: no info for *server:pathname*

Autofs failed to get *pathconf* information for the path name (see the *fpathconf(2)* man page).

pathconf: *server*: server not responding

Autofs is unable to contact the mount daemon on *server* that provides the information to *pathconf()*.

Other Errors With Autofs

If the */etc/auto** files have the execute bit set, the automounter tries to execute the maps, which creates messages such as the following :

```
/etc/auto_home: +auto_home: not found
```

In this situation, the *auto_home* file has incorrect permissions. Each entry in the file generates an error message that is similar to this message. The permissions to the file should be reset by typing the following command:

```
# chmod 644 /etc/auto_home
```

NFS Error Messages

This section shows an error message that is followed by a description of the conditions that should create the error and at minimum one remedy.

Bad argument specified with index option - must be a file
You must include a file name with the index option. You cannot use directory names.

Cannot establish NFS service over /dev/tcp: transport setup problem

This message is often created when the services information in the namespace has not been updated. The message can also be reported for UDP. To fix this problem, you must update the services data in the namespace. For NIS+, the entries should be as follows:

```
nfsd nfsd tcp 2049 NFS server daemon
nfsd nfsd udp 2049 NFS server daemon
```

For NIS and /etc/services, the entries should be as follows:

```
nfsd      2049/tcp      nfs      # NFS server daemon
nfsd      2049/udp      nfs      # NFS server daemon
```

Cannot use index option without public option

Include the public option with the index option in the share command. You must define the public file handle in order for the index option to work.

Note – The Solaris 2.5.1 release required that the public file handle be set by using the share command. A change in the Solaris 2.6 release sets the public file handle to be root (/) by default. This error message is no longer relevant.

Could not start *daemon*: *error*

This message is displayed if the daemon terminates abnormally or if a system call error occurs. The *error* string defines the problem.

Could not use public filehandle in request to *server*

This message is displayed if the public option is specified but the NFS server does not support the public file handle. In this situation, the mount fails. To remedy this situation, either try the mount request without using the public file handle or reconfigure the NFS server to support the public file handle.

daemon running already with pid *pid*

The daemon is already running. If you want to run a new copy, kill the current version and start a new version.

error locking *lock file*

This message is displayed when the *lock file* that is associated with a daemon cannot be locked properly.

error checking *lock file*: *error*

This message is displayed when the *lock file* that is associated with a daemon cannot be opened properly.

NOTICE: NFS3: failing over from *host1* to *host2*

This message is displayed on the console when a failover occurs. The message is advisory only.

filename: File too large

An NFS version 2 client is trying to access a file that is over 2 Gbytes.

mount: ... server not responding:RPC_PMAP_FAILURE - RPC_TIMED_OUT

The server that is sharing the file system you are trying to mount is down or unreachable, at the wrong run level, or its rpcbind is dead or hung.

mount: ... server not responding: RPC_PROG_NOT_REGISTERED

The mount request registered with rpcbind, but the NFS mount daemon mountd is not registered.

mount: ... No such file or directory

Either the remote directory or the local directory does not exist. Check the spelling of the directory names. Run `ls` on both directories.

mount: ...: Permission denied

Your computer name might not be in the list of clients or netgroup that is allowed access to the file system you tried to mount. Use `showmount -e` to verify the access list.

nfs mount: ignoring invalid option "*-option*"

The *-option* flag is not valid. Refer to the `mount_nfs(1M)` man page to verify the required syntax.

Note – This error message is not displayed when running any version of the `mount` command that is included in a Solaris release from 2.6 to the current release or in earlier versions that have been patched.

nfs mount: NFS can't support "*nolargefiles*"

An NFS client has attempted to mount a file system from an NFS server by using the *-nolargefiles* option. This option is not supported for NFS file system types.

nfs mount: NFS V2 can't support "*largefiles*"

The NFS version 2 protocol cannot handle large files. You must use version 3 if access to large files is required.

NFS server *hostname* not responding still trying

If programs hang while doing file-related work, your NFS server might have failed. This message indicates that NFS server *hostname* is down or that a problem has occurred with the server or the network. If failover is being used, *hostname* is a list of servers. Start troubleshooting with "How to Check Connectivity on an NFS Client" on page 183.

NFS fsstat failed for server *hostname*: RPC: Authentication error

This error can be caused by many situations. One of the most difficult situations to debug is when this problem occurs because a user is in too many groups. Currently, a user can be in up to 16 groups but no more if the user is accessing files through NFS mounts. An alternate does exist for users who need to be in more than 16 groups. You can use access control lists to provide the needed access privileges, if you run at minimum the Solaris 2.5 release on the NFS server and the NFS clients.

port *number* in nfs URL not the same as port *number* in port option

The port number that is included in the NFS URL must match the port number that is included with the `-port` option to mount. If the port numbers do not match, the mount fails. Either change the command to make the port numbers identical or do not specify the port number that is incorrect. Usually, you do not need to specify the port number in the NFS URL and with the `-port` option.

replicas must have the same version

For NFS failover to function properly, the NFS servers that are replicas must support the same version of the NFS protocol. Mixing version 2 and version 3 servers is not allowed.

replicated mounts must be read-only

NFS failover does not work on file systems that are mounted read-write. Mounting the file system read-write increases the likelihood that a file could change. NFS failover depends on the file systems being identical.

replicated mounts must not be soft

Replicated mounts require that you wait for a timeout before failover occurs. The `soft` option requires that the mount fail immediately when a timeout starts, so you cannot include the `-soft` option with a replicated mount.

share_nfs: Cannot share more than one filesystem with 'public' option

Check that the `/etc/dfs/dfstab` file has only one file system selected to be shared with the `-public` option. Only one public file handle can be established per server, so only one file system per server can be shared with this option.

Composed September 5, 2002

WARNING: No network locking on *hostname:path*: contact admin to
install server change

An NFS client has unsuccessfully attempted to establish a connection with the network lock manager on an NFS server. Rather than fail the mount, this warning is generated to warn you that locking does not work.

Composed September 5, 2002

Accessing Remote File Systems (Reference)

This chapter provides an introduction to the NFS commands. This chapter also provides information about all of the pieces of the NFS environment and how these pieces work together.

- “NFS Files” on page 197
- “NFS Daemons” on page 200
- “NFS Commands” on page 204
- “Other Useful Commands” on page 217
- “How the NFS Service Works” on page 222
- “Autofs Maps” on page 232
- “How Autofs Works” on page 238
- “Autofs Reference” on page 249

NFS Files

You need several files to support NFS activities on any computer. Many of these files are ASCII, but some of the files are data files. Table 16–1 lists these files and their functions.

TABLE 16–1 NFS Files

File Name	Function
/etc/default/fs	Lists the default file system type for local file systems.
/etc/default/nfs	Lists configuration information for <code>lockd</code> and <code>nfsd</code> .
/etc/default/nfslogd	Lists configuration information for the NFS server logging daemon, <code>nfslogd</code> .

TABLE 16-1 NFS Files (Continued)

File Name	Function
<code>/etc/dfs/dfstab</code>	Lists the local resources to be shared.
<code>/etc/dfs/fstypes</code>	Lists the default file-system types for remote file systems.
<code>/etc/dfs/sharetab</code>	Lists the local and remote resources that are shared. See the <code>sharetab(4)</code> man page. Do not edit this file.
<code>/etc/mnttab</code>	Lists file systems that are currently mounted, including automounted directories. See the <code>mnttab(4)</code> man page. Do not edit this file.
<code>/etc/netconfig</code>	Lists the transport protocols. Do not edit this file.
<code>/etc/nfs/nfslog.conf</code>	Lists general configuration information for NFS server logging.
<code>/etc/nfs/nfslogtab</code>	Lists information for log post-processing by <code>nfslogd</code> . Do not edit this file.
<code>/etc/nfssec.conf</code>	Lists NFS security services. Do not edit this file.
<code>/etc/rmtab</code>	Lists file systems that are remotely mounted by NFS clients. See the <code>rmtab(4)</code> man page. Do not edit this file.
<code>/etc/vfstab</code>	Defines file systems to be mounted locally. See the <code>vfstab(4)</code> man page.

The first entry in `/etc/dfs/fstypes` is often used as the default file-system type for remote file systems. This entry defines the NFS file-system type as the default.

Only one entry is in `/etc/default/fs`: the default file-system type for local disks. You can determine the file-system types that are supported on a client or server by checking the files in `/kernel/fs`.

`/etc/default/nfslogd`

This file defines some of the parameters that are used when using NFS server logging. The following parameters can be defined.

CYCLE_FREQUENCY

Determines the number of hours that must pass before the log files are cycled. The default value is 24 hours. This option is used to prevent the log files from growing too large.

IDLE_TIME

Sets the number of seconds `nfslogd` should sleep before checking for more information in the buffer file. This parameter also determines how often the configuration file is checked. This parameter, along with `MIN_PROCESSING_SIZE`,

determines how often the buffer file is processed. The default value is 300 seconds. Increasing this number can improve performance by reducing the number of checks.

MAPPING_UPDATE_INTERVAL

Specifies the number of seconds between updates of the records in the file-handle-to-path mapping tables. The default value is 86400 seconds or one day. This parameter helps keep the file-handle-to-path mapping tables up-to-date without having to continually update the tables.

MAX_LOGS_PRESERVE

Determines the number of log files to be saved. The default value is 10.

MIN_PROCESSING_SIZE

Sets the minimum number of bytes that the buffer file must reach before processing and writing to the log file. This parameter, along with `IDLE_TIME`, determines how often the buffer file is processed. The default value for is 524288 bytes. Increasing this number can improve performance by reducing the number of times the buffer file is processed.

PRUNE_TIMEOUT

Selects the number of hours that must pass before a file-handle-to-path mapping record times out and can be pruned. The default value is 168 hours or 7 days.

UMASK

Specifies the file mode creation mask for the log files that are created by `nfslogd`. The default value is 0137.

/etc/nfs/nfslog.conf

This file defines the path, file names, and type of logging to be used by `nfslogd`. Each definition is associated with a *tag*. Starting NFS server logging requires that you identify the *tag* for each file system. The global tag defines the default values. You can use the following parameters with each tag as needed.

`defaultdir=`*path*

Specifies the default directory path for the logging files.

`log=`*path/filename*

Sets the path and file name for the log files.

`fhfile=`*path/filename*

Selects the path and file name for the file-handle-to-path database files.

`buffer=`*path/filename*

Determines the path and file name for the buffer files.

`logformat=`*basic | extended*

Selects the format to be used when creating user-readable log files. The basic format produces a log file that is similar to some `ftpd` daemons. The extended format gives a more detailed view.

If the path is not specified, the path that is defined by `defaultdir` is used. Also, you can override `defaultdir` by using an absolute path.

To identify the files more easily, place the files in separate directories. Here is an example of the changes that are needed.

```
% cat /etc/nfs/nfslog.conf
#ident    "@(#)nfslog.conf        1.5        99/02/21  SMI"
#
.
.
# NFS server log configuration file.
#

global    defaultdir=/var/nfs \
          log=nfslog fhtable=fhtable buffer=nfslog_workbuffer

publicftp log=logs/nfslog fhtable=fh/fhtables buffer=buffers/workbuffer
```

In this example, any file system that is shared with `log=publicftp` would use the following values:

- The default directory would be `/var/nfs`.
- Log files would be stored in `/var/nfs/logs/nfslog*`.
- File-handle-to-path database tables would be stored in `/var/nfs/fh/fhtables`.
- Buffer files would be stored in `/var/nfs/buffers/workbuffer`.

NFS Daemons

To support NFS activities, several daemons are started when a system goes into run level 3 or multiuser mode. The `mountd` and `nfsd` daemons are run on systems that are NFS servers. The automatic startup of the server daemons depends on the existence of entries that are labeled with the NFS file-system type in `/etc/dfs/sharetab`. The `lockd` and `statd` daemons are run on NFS clients and servers to support NFS file locking.

automountd

This daemon handles the mount and unmount requests from the `autofs` service. The syntax of the command is as follows:

```
automountd [ -Tnv ] [ -D name=value ]
```

The command behaves in the following ways:

- `-T` enables tracing.

- `-n` disables browsing on all autofs node.
- `-v` selects to log all status messages to the console.
- `-D name=value` substitutes *value* for the automount map variable that is indicated by *name*.

The default value for the automount map is `/etc/auto_master`. Use the `-T` option for troubleshooting.

lockd

This daemon supports record-locking operations on NFS files. The `lockd` daemon manages RPC connections between the client and the server for the Network Lock Manager (NLM) protocol. The daemon is normally started without any options. You can use three options with this command. See the `lockd(1M)` man page. These options can either be used from the command line or by editing the appropriate string in `/etc/default/nfs`. Changing `/etc/default/nfs` makes the change persist through system reboots. This feature is only available in the Solaris 9 release. The only way to make these changes permanent in other releases is to change `/etc/init.d/nfs.client`.

The `LOCKD_GRACE_PERIOD=graceperiod` parameter in `/etc/default/nfs` selects the number of seconds that the clients have to reclaim locks after a server reboot. During this time, the NFS server only processes reclaims of old locks. All other requests for service must wait until the grace period is over. This option affects the NFS server-side response, so this response can be changed only on an NFS server. The default value for *graceperiod* is 45 seconds. Reducing this value means that NFS clients can resume operation more quickly after a server reboot. However, a reduction increases the possibility that a client might not be able to recover all its locks. This same behavior can be used on the command line by starting the daemon with the `-g graceperiod` option.

The `LOCKD_RETRANSMIT_TIMEOUT=timeout` parameter in `/etc/default/nfs` selects the number of seconds to wait before retransmitting a lock request to the remote server. This option affects the NFS client-side service. The default value for *timeout* is 15 seconds. Decreasing the *timeout* value can improve response time for NFS clients on a noisy network. However, this change can cause additional server load by increasing the frequency of lock requests. The same parameter can be used from the command line by starting the daemon with the `-t timeout` option.

The `LOCKD_SERVERS=nthreads` parameter in `/etc/default/nfs` specifies the maximum number of concurrent threads that the server handles per connection. Base the value for *nthreads* on the load that is expected on the NFS server. The default value is 20. Each NFS client that uses TCP uses a single connection with the NFS server. Therefore, each client can use a maximum of 20 concurrent threads on the server. All NFS clients that use UDP share a single connection with the NFS server. Under these conditions, you might have to increase the number of threads that are available for the

UDP connection. A minimum calculation would be to allow two threads for each UDP client. However, this number is specific to the workload on the client, so two threads per client might not be sufficient. The disadvantage to using more threads is that when the threads are used, more memory is used on the NFS server. If the threads are never used, however, increasing *nthreads* has no effect. The same parameter can be used from the command line by starting the daemon with the *nthreads* option.

mountd

This daemon handles file-system mount requests from remote systems and provides access control. The *mountd* daemon checks */etc/dfs/sharetab* to determine which file systems are available for remote mounting and which systems are allowed to do the remote mounting. You can use the *-v* option and the *-r* option with this command. See the *mountd(1M)* man page.

The *-v* option runs the command in verbose mode. Every time an NFS server determines the access that a client should get, a message is printed on the console. The information that is generated can be useful when trying to determine why a client cannot access a file system.

The *-r* option rejects all future mount requests from clients. This option does not affect clients that already have a file system mounted.

nfsd

This daemon handles other client file-system requests. You can use several options with this command. See the *nfsd(1M)* man page for a complete listing. These options can either be used from the command line or by editing the appropriate string in */etc/default/nfs*. Changing */etc/default/nfs* makes the change persist through system reboots. This feature is only available in the Solaris 9 release. The only way to make these changes permanent in other releases is to change */etc/init.d/nfs.server*.

The *NFSD_LISTEN_BACKLOG=length* parameter in */etc/default/nfs* sets the length of the connection queue over connection-oriented transports for NFS and TCP. The default value is 32 entries. The same selection can be made from the command line by starting *nfsd* with the *-l* option.

The *NFSD_MAX_CONNECTIONS=#_conn* parameter in */etc/default/nfs* selects the maximum number of connections per connection-oriented transport. The default value for *#_conn* is unlimited. The same parameter can be used from the command line by starting the daemon with the *-c #_conn* option.

The *NFSD_SERVER=nservers* parameter in */etc/default/nfs* selects the maximum number of concurrent requests that a server can handle. The default value for *nservers* is 1, but the startup scripts select 16. The same selection can be made from the command line by starting *nfsd* with the *nservers* option.

Unlike older versions of this daemon, `nfsd` does not spawn multiple copies to handle concurrent requests. Checking the process table with `ps` only shows one copy of the daemon running.

`nfslogd`

This daemon provides operational logging. NFS operations that are logged against a server are based on the configuration options that are defined in `/etc/default/nfslogd`. When NFS server logging is enabled, records of all RPC operations on a selected file system are written to a buffer file by the kernel. Then `nfslogd` post-processes these requests. The name service switch is used to help map UIDs to logins and IP addresses to host names. The number is recorded if no match can be found through the identified name services.

Mapping of file handles to path names is also handled by `nfslogd`. The daemon tracks these mappings in a file-handle-to-path mapping table. One mapping table exists for each tag that is identified in `/etc/nfs/nfslogd`. After post-processing, the records are written to ASCII log files.

`statd`

This daemon works with `lockd` to provide crash and recovery functions for the lock manager. The `statd` daemon tracks the clients that hold locks on an NFS server. If a server crashes, on rebooting `statd` on the server contacts `statd` on the client. The client `statd` can then attempt to reclaim any locks on the server. The client `statd` also informs the server `statd` when a client has crashed so that the client's locks on the server can be cleared. You have no options to select with this daemon. For more information, see the `statd(1M)` man page.

In the Solaris 7 release, the way that `statd` tracks the clients has been improved. In all earlier Solaris releases, `statd` created files in `/var/statmon/sm` for each client by using the client's unqualified host name. This file naming caused problems if you had two clients in different domains that shared a host name, or if clients were not resident in the same domain as the NFS server. Because the unqualified host name only lists the host name, without any domain or IP-address information, the older version of `statd` had no way to differentiate between these types of clients. To fix this problem, the Solaris 7 `statd` creates a symbolic link in `/var/statmon/sm` to the unqualified host name by using the IP address of the client. The new link resembles the following:

```
# ls -l /var/statmon/sm
lrwxrwxrwx 1 daemon 11 Apr 29 16:32 ipv4.192.9.200.1 -> myhost
lrwxrwxrwx 1 daemon 11 Apr 29 16:32 ipv6.fec0::56:a00:20ff:feb9:2734 -> v6host
--w----- 1 daemon 11 Apr 29 16:32 myhost
--w----- 1 daemon 11 Apr 29 16:32 v6host
```

In this example, the client host name is `myhost` and the client's IP address is `192.9.200.1`. If another host with the name `myhost` were mounting a file system, two symbolic links would lead to the host name.

NFS Commands

These commands must be run as root to be fully effective, but requests for information can be made by all users:

- “`automount`” on page 204
- “`clear_locks`” on page 205
- “`mount`” on page 205
- “`mountall`” on page 209
- “`setmnt`” on page 216
- “`share`” on page 210
- “`shareall`” on page 215
- “`showmount`” on page 215
- “`umount`” on page 208
- “`umountall`” on page 209
- “`unshare`” on page 214
- “`unshareall`” on page 215

automount

This command installs autofs mount points and associates the information in the automaster files with each mount point. The syntax of the command is as follows:

```
automount [ -t duration ] [ -v ]
```

`-t duration` sets the time, in seconds, that a file system is to remain mounted, and `-v` selects the verbose mode. Running this command in the verbose mode allows for easier troubleshooting.

If not specifically set, the value for duration is set to 5 minutes. In most circumstances, this value is good. However, on systems that have many automounted file systems, you might need to increase the duration value. In particular, if a server has many users active, checking the automounted file systems every 5 minutes can be inefficient. Checking the autofs file systems every 1800 seconds, which is 30 minutes, could be more optimal. By not unmounting the file systems every 5 minutes, `/etc/mnttab` can become large. To reduce the output when `df` checks each entry in `/etc/mnttab`, you can filter the output from `df` by using the `-F` option (see the `df(1M)` man page) or by using `egrep`.

You should consider the idea that adjusting the duration also changes how quickly changes to the automounter maps are reflected. Changes cannot be seen until the file system is unmounted. Refer to “Modifying the Maps” on page 171 for instructions on how to modify automounter maps.

clear_locks

This command enables you to remove all file, record, and share locks for an NFS client. You must be `root` to run this command. From an NFS server, you can clear the locks for a specific client. From an NFS client, you can clear locks for that client on a specific server. The following example would clear the locks for the NFS client that is named `tulip` on the current system.

```
# clear_locks tulip
```

Using the `-s` option enables you to specify which NFS host to clear the locks from. You must run this option from the NFS client, which created the locks. In this situation, the locks from the client would be removed from the NFS server that is named `bee`.

```
# clear_locks -s bee
```



Caution – This command should only be run when a client crashes and cannot clear its locks. To avoid data corruption problems, do not clear locks for an active client.

mount

With this command, you can attach a named file system, either local or remote, to a specified mount point. For more information, see the `mount(1M)` man page. Used without arguments, `mount` displays a list of file systems that are currently mounted on your computer.

Many types of file systems are included in the standard Solaris installation. Each file-system type has a specific man page that lists the options to `mount` that are appropriate for that file-system type. The man page for NFS file systems is `mount_nfs(1M)`. For UFS file systems, see `mount_ufs(1M)`.

The Solaris 7 release includes the ability to select a path name to mount from an NFS server by using an NFS URL instead of the standard `server:/pathname` syntax. See “How to Mount an NFS File System Using an NFS URL” on page 161 for further information.



Caution – The version of the `mount` command that is included in any Solaris release from 2.6 to the current release does not warn about invalid options. The command silently ignores any options that cannot be interpreted. Ensure that you verify all of the options that were used so that you can prevent unexpected behavior.

`mount` Options for NFS File Systems

The subsequent text lists some of the options that can follow the `-o` flag when you are mounting an NFS file system.

`bg | fg`

These options can be used to select the retry behavior if a mount fails. The `-bg` option causes the mount attempts to be run in the background. The `-fg` option causes the mount attempt to be run in the foreground. The default is `-fg`, which is the best selection for file systems that must be available. This option prevents further processing until the mount is complete. `-bg` is a good selection for non-critical file systems, because the client can do other processing while waiting for the mount request to complete.

`forcedirectio`

This option improves performance of sequential reads on large files. Data is copied directly to a user buffer. No caching is done in the kernel on the client. This option is off by default.

`largefiles`

With this option, you can access files that are larger than 2 Gbytes on a server that is running the Solaris 2.6 release. Whether a large file can be accessed can only be controlled on the server, so this option is silently ignored on NFS version 3 mounts. Starting with release 2.6, by default, all UFS file systems are mounted with `largefiles`. For mounts that use the NFS version 2 protocol, the `largefiles` option causes the mount to fail with an error.

`nolargefiles`

This option for UFS mounts guarantees that no large files can exist on the file system. See the `mount_ufs(1M)` man page. Because the existence of large files can only be controlled on the NFS server, no option for `nolargefiles` exists when using NFS mounts. Attempts to NFS-mount a file system by using this option are rejected with an error.

`public`

This option forces the use of the public file handle when contacting the NFS server. If the public file handle is supported by the server, the mounting operation is faster because the MOUNT protocol is not used. Also, because the MOUNT protocol is not used, the `public` option allows mounting to occur through a firewall.

`rw | ro`

The `-rw` and `-ro` options indicate whether a file system is to be mounted read-write or read-only. The default is read-write, which is the appropriate option

for remote home directories, mail-spooling directories, or other file systems that need to be changed by users. The read-only option is appropriate for directories that should not be changed by users. For example, shared copies of the man pages should not be writable by users.

sec=mode

You can use this option to specify the authentication mechanism to be used during the mount transaction. The value for *mode* can be one of the values that is shown in Table 16-2. The modes are also defined in `/etc/nfssec.conf`.

TABLE 16-2 NFS Security Modes

Mode	Selected Authentication Service
krb5	Kerberos Version 5
krb5i	Kerberos Version 5 with integrity
krb5p	Kerberos Version 5 with privacy
none	No authentication
dh	Diffie-Hellman (DH) authentication
sys	Standard UNIX authentication

soft | hard

An NFS file system that is mounted with the *soft* option returns an error if the server does not respond. The *hard* option causes the mount to continue to retry until the server responds. The default is *hard*, which should be used for most file systems. Applications frequently do not check return values from *soft*-mounted file systems, which can make the application fail or can lead to corrupted files. If the application does check the return values, routing problems and other conditions can still confuse the application or lead to file corruption if the *soft* option is used. In most situations, the *soft* option should not be used. If a file system is mounted by using the *hard* option and becomes unavailable, an application that uses this file system hangs until the file system becomes available.

Using the mount Command

Both of these commands mount an NFS file system from the server *bee* read-only:

```
# mount -F nfs -r bee:/export/share/man /usr/man
# mount -F nfs -o ro bee:/export/share/man /usr/man
```

This command uses the *-O* option to force the man pages from the server *bee* to be mounted on the local system even if `/usr/man` has already been mounted. See the following:

```
# mount -F nfs -O bee:/export/share/man /usr/man
```

This command uses client failover:

```
# mount -F nfs -r bee,wasp:/export/share/man /usr/man
```

Note – When used from the command line, the listed servers must support the same version of the NFS protocol. Do not mix version 2 and version 3 servers when running mount from the command line. You can use mixed servers with autofs. Autofs automatically selects the best subset of version 2 or version 3 servers.

Here is an example of using an NFS URL with the mount command:

```
# mount -F nfs nfs://bee//export/share/man /usr/man
```

Use the mount command with no arguments to display file systems that are mounted on a client.

```
% mount
/ on /dev/dsk/c0t3d0s0 read/write/setuid on Tues Jan 24 13:20:47 1995
/usr on /dev/dsk/c0t3d0s6 read/write/setuid on Tues Jan 24 13:20:47 1995
/proc on /proc read/write/setuid on Tues Jan 24 13:20:47 1995
/dev/fd on fd read/write/setuid on Tues Jan 24 13:20:47 1995
/tmp on swap read/write on Tues Jan 24 13:20:51 1995
/opt on /dev/dsk/c0t3d0s5 setuid/read/write on Tues Jan 24 13:20:51 1995
/home/kathys on bee:/export/home/bee7/kathys
intr/noquota/nosuid/remote on Tues Jan 24 13:22:13 1995
```

umount

This command enables you to remove a remote file system that is currently mounted. The umount command supports the -V option to allow for testing. You might also use the -a option to unmount several file systems at one time. If *mount_points* are included with the -a option, those file systems are unmounted. If no mount points are included, an attempt is made to unmount all file systems that are listed in */etc/mnttab*, except for the “required” file systems, such as */*, */usr*, */var*, */proc*, */dev/fd*, and */tmp*. Because the file system is already mounted and should have an entry in */etc/mnttab*, you do not need to include a flag for the file-system type.

The -f option forces a busy file system to be unmounted. You can use this option to unhang a client that is hung while trying to mount an unmountable file system.



Caution – By forcing an unmount of a file system, you can cause data loss if files are being written to.

Using the `umount` Command

This example unmounts a file system that is mounted on `/usr/man`:

```
# umount /usr/man
```

This example displays the results of running `umount -a -V`:

```
# umount -a -V
umount /home/kathys
umount /opt
umount /home
umount /net
```

Notice that this command does not actually unmount the file systems.

`mountall`

Use this command to mount all file systems or a specific group of file systems that are listed in a file-system table. The command provides a way to do the following:

- Select the file-system type to be accessed with the `-F FSType` option
- Select all the remote file systems that are listed in a file-system table with the `-r` option
- Select all the local file systems with the `-l` option

Because all file systems that are labeled as NFS file-system type are remote file systems, some of these options are redundant. For more information, see the `mountall(1M)` man page.

Using the `mountall` Command

These two examples are equivalent:

```
# mountall -F nfs
# mountall -F nfs -r
```

`umountall`

Use this command to unmount a group of file systems. The `-k` option runs the `fuser -k mount_point` command to kill any processes that are associated with the *mount_point*. The `-s` option indicates that unmount is not to be performed in parallel. `-l` specifies that only local file systems are to be used, and `-r` specifies that only remote file systems are to be used. The `-h host` option indicates that all file systems from the named host should be unmounted. You cannot combine the `-h` option with `-l` or `-r`.

Using the `umountall` Command

This command unmounts all file systems that are mounted from remote hosts:

```
# umountall -r
```

This command unmounts all file systems that are currently mounted from the server `bee`:

```
# umountall -h bee
```

`share`

With this command, you can make a local file system on an NFS server available for mounting. You can also use the `share` command to display a list of the file systems on your system that are currently shared. The NFS server must be running in order for the `share` command to work. The NFS server software is started automatically during boot if an entry is in `/etc/dfs/dfstab`. The command does not report an error if the NFS server software is not running, so you yourself must verify that the software is running.

The objects that can be shared include any directory tree. However, each file system hierarchy is limited by the disk slice or partition that the file system is located on. For instance, sharing the root (`/`) file system would not also share `/usr`, unless these directories are on the same disk partition or slice. Normal installation places root on slice 0 and `/usr` on slice 6. Also, sharing `/usr` would not share any other local disk partitions that are mounted on subdirectories of `/usr`.

A file system cannot be shared if that file system is part of a larger file system that is already being shared. For example, if `/usr` and `/usr/local` are on one disk slice, `/usr` can be shared or `/usr/local` can be shared. However, if both directories need to be shared with different share options, `/usr/local` must be moved to a separate disk slice.

Note – You can gain access to a file system that is read-only shared through the file handle of a file system that is read-write shared. However, the two file systems have to be on the same disk slice. You can create a more secure situation. Place those file systems that need to be read-write on a separate partition or separate disk slice from the file systems that you need to share as read-only.

Non-File System-Specific share Options

Some of the options that you can include with the `-o` flag are as follows.

`rw | ro`

The *pathname* file system is shared read-write or read-only for all clients.

rw=accesslist

The file system is shared read-write for the clients that are listed only. All other requests are denied. Starting with the Solaris 2.6 release, the list of clients that are defined in *accesslist* has been expanded. See “Setting Access Lists With the *share* Command” on page 213 for more information. You can use this option to override an *-ro* option.

NFS-Specific share Options

The options that you can use with NFS file systems include the following.

aclok

This option enables an NFS server that supports the NFS version 2 protocol to be configured to do access control for NFS version 2 clients. Without this option, all clients are given minimal access. With this option, the clients have maximal access. For instance, on file systems that are shared with the *-aclok* option, if anyone has read permissions, everyone does. However, without this option, you can deny access to a client who should have access permissions. A decision to permit too much access or too little access depends on the security systems already in place. See “Using Access Control Lists (ACLs)” in *System Administration Guide: Security Services* for more information about access control lists (ACLs).

Note – To use ACLs, ensure that clients and servers run software that supports the NFS version 3 and NFS_ACL protocols. If the software only supports the NFS version 3 protocol, clients obtain correct access but cannot manipulate the ACLs. If the software supports the NFS_ACL protocol, the clients obtain correct access and can manipulate the ACLs. Starting with release 2.5, the Solaris system supports both protocols.

anon=uid

You use *uid* to select the user ID of unauthenticated users. If you set *uid* to *-1*, the server denies access to unauthenticated users. You can grant root access by setting *anon=0*, but this option allows unauthenticated users to have root access, so use the *root* option instead.

index=filename

When a user accesses an NFS URL, the *-index=filename* option forces the HTML file to load, instead of displaying a list of the directory. This option mimics the action of current browsers if an *index.html* file is found in the directory that the HTTP URL is accessing. This option is the equivalent of setting the *DirectoryIndex* option for *httpd*. For instance, suppose that the *dfstab* file entry resembles the following:

```
share -F nfs -o ro,public,index=index.html /export/web
```

These URLs then display the same information:

```
nfs://<server>/<dir>  
nfs://<server>/<dir>/index.html  
nfs://<server>/export/web/<dir>  
nfs://<server>/export/web/<dir>/index.html  
http://<server>/<dir>  
http://<server>/<dir>/index.html
```

log=tag

This option specifies the tag in `/etc/nfs/nfslog.conf` that contains the NFS server logging configuration information for a file system. This option must be selected to enable NFS server logging.

nosuid

This option signals that all attempts to enable the `setuid` or `setgid` mode should be ignored. NFS clients cannot create files with the `setuid` or `setgid` bits on.

public

The `-public` option has been added to the `share` command to enable WebNFS browsing. Only one file system on a server can be shared with this option.

root=accesslist

The server gives root access to the hosts in the list. By default, the server does not give root access to any remote hosts. If the selected security mode is anything other than `-sec=sys`, you can only include client host names in the *accesslist*. Starting with the Solaris 2.6 release, the list of clients that are defined in *accesslist* is expanded. See "Setting Access Lists With the `share` Command" on page 213 for more information.



Caution – Granting root access to other hosts has wide security implications. Use the `-root=` option with extreme caution.

sec=mode[:mode]

mode selects the security modes that are needed to obtain access to the file system. By default, the security mode is UNIX authentication. You can specify multiple modes, but use each security mode only once per command line. Each `-mode` option applies to any subsequent `-rw`, `-ro`, `-rw=`, `-ro=`, `-root=`, and `-window=` options until another `-mode` is encountered. The use of `-sec=none` maps all users to user `nobody`.

window=value

value selects the maximum lifetime in seconds of a credential on the NFS server. The default value is 30000 seconds or 8.3 hours.

Setting Access Lists With the share Command

In Solaris releases prior to 2.6, the *accesslist* that was included with either the `-ro=`, `-rw=`, or `-root=` option of the `share` command was restricted to a list of host names or netgroup names. Starting with the Solaris 2.6 release, the access list can also include a domain name, a subnet number, or an entry to deny access. These extensions should simplify file access control on a single server without having to change the namespace or maintain long lists of clients.

This command provides read-only access for most systems but allows read-write access for `rose` and `lilac`:

```
# share -F nfs -o ro,rw=rose:lilac /usr/src
```

In the next example, read-only access is assigned to any host in the `eng` netgroup. The client `rose` is specifically given read-write access.

```
# share -F nfs -o ro=eng,rw=rose /usr/src
```

Note – You cannot specify both `rw` and `ro` without arguments. If no read-write option is specified, the default is read-write for all clients.

To share one file system with multiple clients, you must type all options on the same line. Multiple invocations of the `share` command on the same object “remember” only the last command that is run. This command enables read-write access to three client systems, but only `rose` and `tulip` are given access to the file system as `root`.

```
# share -F nfs -o rw=rose:lilac:tulip,root=rose:tulip /usr/src
```

When sharing a file system that uses multiple authentication mechanisms, ensure that you include the `-ro`, `-ro=`, `-rw`, `-rw=`, `-root`, and `-window` options after the correct security modes. In this example, UNIX authentication is selected for all hosts in the netgroup that is named `eng`. These hosts can only mount the file system in read-only mode. The hosts `tulip` and `lilac` can mount the file system read-write if these hosts use Diffie-Hellman authentication. With these options, `tulip` and `lilac` can mount the file system read-only even if these hosts are not using DH authentication. However, the host names must be listed in the `eng` netgroup.

```
# share -F nfs -o sec=dh,rw=tulip:lilac,sec=sys,ro=eng /usr/src
```

Even though UNIX authentication is the default security mode, UNIX authentication is not included if the `-sec` option is used. Therefore, you must include a `-sec=sys` option if UNIX authentication is to be used with any other authentication mechanism.

You can use a DNS domain name in the access list by preceding the actual domain name with a dot. The string that follows the dot is a domain name, not a fully qualified host name. The following entry allows mount access to all hosts in the `eng.example.com` domain:

```
# share -F nfs -o ro=.:eng.example.com /export/share/man
```

In this example, the single “.” matches all hosts that are matched through the NIS or NIS+ namespaces. The results that are returned from these name services do not include the domain name. The “.eng.example.com” entry matches all hosts that use DNS for namespace resolution. DNS always returns a fully qualified host name. So, the longer entry is required if you use a combination of DNS and the other namespaces.

You can use a subnet number in an access list by preceding the actual network number or the network name with “@”. This character differentiates the network name from a netgroup or a fully qualified host name. You must identify the subnet in either /etc/networks or in an NIS or NIS+ namespace. The following entries have the same effect if the 129.144 subnet has been identified as the eng network:

```
# share -F nfs -o ro=@eng /export/share/man
# share -F nfs -o ro=@129.144 /export/share/man
# share -F nfs -o ro=@129.144.0.0 /export/share/man
```

The last two entries show that you do not need to include the full network address.

If the network prefix is not byte aligned, as with Classless Inter-Domain Routing (CIDR), the mask length can be explicitly specified on the command line. The mask length is defined by following either the network name or the network number with a slash and the number of significant bits in the prefix of the address. For example:

```
# share -f nfs -o ro=@eng/17 /export/share/man
# share -F nfs -o ro=@129.144.132/17 /export/share/man
```

In these examples, the “/17” indicates that the first 17 bits in the address are to be used as the mask. For additional information on CIDR, look up RFC 1519.

You can also select negative access by placing a “-” before the entry. Note that the entries are read from left to right. Therefore, you must place the negative access entries before the entry that the negative access entries apply to:

```
# share -F nfs -o ro=-rose.:eng.example.com /export/share/man
```

This example would allow access to any hosts in the eng.example.com domain except the host that is named rose.

unshare

This command allows you to make a previously available file system unavailable for mounting by clients. You can use the unshare command to unshare any file system—whether the file system was shared explicitly with the share command or automatically through /etc/dfs/dfstab. If you use the unshare command to unshare a file system that you shared through the dfstab file, be careful. Remember that the file system is shared again when you exit and reenter run level 3. You must remove the entry for this file system from the dfstab file if the change is to continue.

When you unshare an NFS file system, access from clients with existing mounts is inhibited. The file system might still be mounted on the client, but the files are not accessible.

Using the unshare Command

This command unshares a specific file system:

```
# unshare /usr/src
```

shareall

This command allows for multiple file systems to be shared. When used with no options, the command shares all entries in `/etc/dfs/dfstab`. You can include a file name to specify the name of a file that lists share command lines. If you do not include a file name, `/etc/dfs/dfstab` is checked. If you use a `"-"` to replace the file name, you can type share commands from standard input.

Using the shareall Command

This command shares all file systems that are listed in a local file:

```
# shareall /etc/dfs/special_dfstab
```

unshareall

This command makes all currently shared resources unavailable. The `-F FSType` option selects a list of file-system types that are defined in `/etc/dfs/fstypes`. This flag enables you to choose only certain types of file systems to be unshared. The default file system type is defined in `/etc/dfs/fstypes`. To choose specific file systems, use the unshare command.

Using the unshareall Command

This example should unshare all NFS-type file systems:

```
# unshareall -F nfs
```

showmount

This command displays one of the following:

- All clients that have remotely mounted file systems that are shared from an NFS server
- Only the file systems that are mounted by clients
- The shared file systems with the client access information

The command syntax is as follows:

```
showmount [ -ade ] [ hostname ]
```

-a	Prints a list of all the remote mounts. Each entry includes the client name and the directory.
-d	Prints a list of the directories that are remotely mounted by clients.
-e	Prints a list of the files that are shared or are exported.
<i>hostname</i>	Selects the NFS server to gather the information from.

If *hostname* is not specified the local host is queried.

Using the showmount Command

This command lists all clients and the local directories that the clients have mounted.

```
# showmount -a bee
lilac:/export/share/man
lilac:/usr/src
rose:/usr/src
tulip:/export/share/man
```

This command lists the directories that have been mounted.

```
# showmount -d bee
/export/share/man
/usr/src
```

This command lists file systems that have been shared.

```
# showmount -e bee
/usr/src                               (everyone)
/export/share/man                     eng
```

setmnt

This command creates an `/etc/mnttab` table. The `mount` and `umount` commands consult the table. Generally, you do not have to run this command manually, as this command runs automatically when a system is booted.

Other Useful Commands

These commands can be useful when troubleshooting NFS problems.

`nfsstat`

You can use this command to gather statistical information about NFS and RPC connections. The syntax of the command is as follows:

```
nfsstat [ -cmnrzs ]
```

-c	Displays client-side information
-m	Displays statistics for each NFS-mounted file system
-n	Specifies that NFS information is to be displayed on both the client side and the server side
-r	Displays RPC statistics
-s	Displays the server-side information
-z	Specifies that the statistics should be set to zero

If no options are supplied on the command line, the `-cnrs` options are used.

Gathering server-side statistics can be important for debugging problems when new software or new hardware is added to the computing environment. Running this command a minimum of once a week, and storing the numbers, provides a good history of previous performance.

Using the `nfsstat` Command

```
# nfsstat -s
```

```
Server rpc:
```

```
Connection oriented:
```

calls	badcalls	nullrecv	badlen	xdrCALL	dupchecks	dupreqs
11420263	0	0	0	0	1428274	19

```
Connectionless:
```

calls	badcalls	nullrecv	badlen	xdrCALL	dupchecks	dupreqs
14569706	0	0	0	0	953332	1601

```
Server nfs:
```

```
calls      badcalls
```

Composed September 5, 2002

```
24234967 226
Version 2: (13073528 calls)
null      getattr  setattr  root      lookup    readlink  read
138612 1% 1192059 9% 45676 0% 0 0%      9300029 71% 9872 0% 1319897 10%
wrcache   write     create   remove    rename    link      symlink
0 0%      805444 6% 43417 0% 44951 0% 3831 0% 4758 0% 1490 0%
mkdir     rmdir    readdir  statfs
2235 0%   1518 0% 51897 0% 107842 0%
Version 3: (11114810 calls)
null      getattr  setattr  lookup    access    readlink  read
141059 1% 3911728 35% 181185 1% 3395029 30% 1097018 9% 4777 0% 960503 8%
write     create   mkdir    symlink    mknod     remove    rmdir
763996 6% 159257 1% 3997 0% 10532 0% 26 0% 164698 1% 2251 0%
rename    link     readdir  readdirplus fsstat    fsinfo    pathconf
53303 0% 9500 0% 62022 0% 79512 0% 3442 0% 34275 0% 3023 0%
commit
73677 0%

Server nfs_acl:
Version 2: (1579 calls)
null      getacl    setacl    getattr    access
0 0%      3 0%      0 0%      1000 63% 576 36%
Version 3: (45318 calls)
null      getacl    setacl
0 0%      45318 100% 0 0%
```

The previous listing is an example of NFS server statistics. The first five lines relate to RPC and the remaining lines report NFS activities. In both sets of statistics, knowing the average number of badcalls or calls and the number of calls per week can help identify a problem. The badcalls value reports the number of bad messages from a client. This value can point out network hardware problems.

Some of the connections generate write activity on the disks. A sudden increase in these statistics could indicate trouble and should be investigated. For NFS version 2 statistics, the connections to note are setattr, write, create, remove, rename, link, symlink, mkdir, and rmdir. For NFS version 3 statistics, the value to watch is commit. If the commit level is high in one NFS server as compared to another almost identical server, check that the NFS clients have enough memory. The number of commit operations on the server grows when clients do not have available resources.

pstack

This command displays a stack trace for each process. The pstack command must be run by the owner of the process or by root. You can use pstack to determine where a process is hung. The only option that is allowed with this command is the PID of the process that you want to check. See the proc(1) man page.

The following example is checking the nfsd process that is running.

```
# /usr/proc/bin/pstack 243
243: /usr/lib/nfs/nfsd -a 16
ef675c04 poll (24d50, 2, ffffffff)
```

```
000115dc ???????? (24000, 132c4, 276d8, 1329c, 276d8, 0)
00011390 main      (3, effffff14, 0, 0, ffffffff, 400) + 3c8
00010fb0 _start    (0, 0, 0, 0, 0, 0) + 5c
```

The example shows that the process is waiting for a new connection request, which is a normal response. If the stack shows that the process is still in poll after a request is made, the process might be hung. Follow the instructions in “How to Restart NFS Services” on page 187 to fix this problem. Review the instructions in “NFS Troubleshooting Procedures” on page 183 to fully verify that your problem is a hung program.

rpcinfo

This command generates information about the RPC service that is running on a system. You can also use this command to change the RPC service. Many options are available with this command. See the `rpcinfo(1M)` man page. The following is a shortened synopsis for some of the options that you can use with the command:

```
rpcinfo [ -m | -s ] [ hostname ]
```

```
rpcinfo -T transport hostname [ progname ]
```

```
rpcinfo [ -t | -u ] [ hostname ] [ progname ]
```

<code>-m</code>	Displays a table of statistics of the <code>rpcbind</code> operations
<code>-s</code>	Displays a concise list of all registered RPC programs
<code>-T</code>	Displays information about services that use specific transports or protocols
<code>-t</code>	Probes the RPC programs that use TCP
<code>-u</code>	Probes the RPC programs that use UDP
<i>transport</i>	Selects the transport or protocol for the services
<i>hostname</i>	Selects the host name of the server you need information from
<i>progname</i>	Selects the RPC program to gather information about

If no value is given for *hostname*, the local host name is used. You can substitute the RPC program number for *progname*, but many users can remember the name and not the number. You can use the `-p` option in place of the `-s` option on those systems that do not run the NFS version 3 software.

The data that is generated by this command can include the following:

- The RPC program number
- The version number for a specific program
- The transport protocol that is being used
- The name of the RPC service
- The owner of the RPC service

Using the rpcinfo Command

This example gathers information on the RPC services that are running on a server. The text that is generated by the command is filtered by the `sort` command to make the output more readable. Several lines that list RPC services have been deleted from the example.

```
% rpcinfo -s bee | sort -n
program version(s) netid(s) service owner
100000 2,3,4 udp6,tcp6,udp,tcp,ticlts,ticotsord,ticots rpcbind superuser
100001 4,3,2 ticlts,udp,udp6 rstatd superuser
100002 3,2 ticots,ticotsord,tcp,tcp6,ticlts,udp,udp6 rusersd superuser
100003 3,2 tcp,udp,tcp6,udp6 nfs superuser
100005 3,2,1 ticots,ticotsord,tcp,tcp6,ticlts,udp,udp6 mountd superuser
100007 1,2,3 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 ypbind superuser
100008 1 ticlts,udp,udp6 walld superuser
100011 1 ticlts,udp,udp6 rquotad superuser
100012 1 ticlts,udp,udp6 sprayd superuser
100021 4,3,2,1 tcp,udp,tcp6,udp6 nlockmgr superuser
100024 1 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 status superuser
100029 3,2,1 ticots,ticotsord,ticlts keyserd superuser
100068 5 tcp,udp cmsd superuser
100083 1 tcp,tcp6 ttldbserverd superuser
100099 3 ticotsord autofs superuser
100133 1 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 - superuser
100134 1 ticotsord tokenring superuser
100155 1 ticots,ticotsord,tcp,tcp6 smserverd superuser
100221 1 tcp,tcp6 - superuser
100227 3,2 tcp,udp,tcp6,udp6 nfs_acl superuser
100229 1 tcp,tcp6 metad superuser
100230 1 tcp,tcp6 metamhd superuser
100231 1 ticots,ticotsord,ticlts - superuser
100232 10 udp,udp6 sadmind superuser
100234 1 ticotsord gssd superuser
100235 1 tcp,tcp6 - superuser
100242 1 tcp,tcp6 metamedd superuser
100249 1 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 - superuser
300326 4 tcp,tcp6 - superuser
300598 1 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 - superuser
390113 1 tcp - unknown
805306368 1 ticots,ticotsord,ticlts,tcp,udp,tcp6,udp6 - superuser
1289637086 1,5 tcp - 26069
```

This example shows how to gather information about a particular RPC service by selecting a particular transport on a server.

```
% rpcinfo -t bee mountd
program 100005 version 1 ready and waiting
program 100005 version 2 ready and waiting
program 100005 version 3 ready and waiting
% rpcinfo -u bee nfs
program 100003 version 2 ready and waiting
program 100003 version 3 ready and waiting
```

The first example checks the `mountd` service that is running over TCP. The second example checks the NFS service that is running over UDP.

snoop

This command is often used to watch for packets on the network. The `snoop` command must be run as `root`. The use of this command is a good way to ensure that the network hardware is functioning on both the client and the server. Many options are available. See the `snoop(1M)` man page. A shortened synopsis of the command follows:

```
snoop [ -d device ] [ -o filename ] [ host hostname ]
```

<code>-d <i>device</i></code>	Specifies the local network interface
<code>-o <i>filename</i></code>	Stores all the captured packets into the named file
<code><i>hostname</i></code>	Displays packets going to and from a specific host only

The `-d device` option is useful on those servers that have multiple network interfaces. You can use many other expressions besides setting the host. A combination of command expressions with `grep` can often generate data that is specific enough to be useful.

When troubleshooting, make sure that packets are going to and from the proper host. Also, look for error messages. Saving the packets to a file can simplify the review of the data.

truss

You can use this command to see if a process is hung. The `truss` command must be run by the owner of the process or by `root`. You can use many options with this command. See the `truss(1)` man page. A shortened syntax of the command follows:

```
truss [ -t syscall ] -p pid
```

<code>-t syscall</code>	Selects system calls to trace
<code>-p pid</code>	Indicates the PID of the process to be traced

The *syscall* can be a comma-separated list of system calls to be traced. Also, starting *syscall* with an `!` selects to exclude the listed system calls from the trace.

This example shows that the process is waiting for another connection request from a new client.

```
# /usr/bin/truss -p 243
poll(0x00024D50, 2, -1)      (sleeping...)
```

The previous example shows a normal response. If the response does not change after a new connection request has been made, the process could be hung. Follow the instructions in “How to Restart NFS Services” on page 187 to fix the hung program. Review the instructions in “NFS Troubleshooting Procedures” on page 183 to fully verify that your problem is a hung program.

How the NFS Service Works

The following sections describe some of the complex functions of the NFS software.

Version 2 and Version 3 Negotiation

NFS servers might be supporting clients that are not using the NFS version 3 software. So, part of the initiation procedure includes negotiation of the protocol level. If both the client and the server can support version 3, that version is used. If either the client or the server can only support version 2, that version is selected.

You can override the values that are determined by the negotiation by using the `-vers` option with the `mount` command. See the `mount_nfs(1M)` man page. Under most circumstances, you should not have to specify the version level, as the best level is selected by default.

UDP and TCP Negotiation

During initiation, the transport protocol is also negotiated. By default, the first connection-oriented transport that is supported on both the client and the server is selected. If this selection does not succeed, the first available connectionless transport protocol is used. The transport protocols that are supported on a system are listed in `/etc/netconfig`. TCP is the connection-oriented transport protocol that is supported by the release. UDP is the connectionless transport protocol.

When both the NFS protocol version and the transport protocol are determined by negotiation, the NFS protocol version is given precedence over the transport protocol. The NFS version 3 protocol that uses UDP is given higher precedence than the NFS version 2 protocol that is using TCP. You can manually select both the NFS protocol version and the transport protocol with the `mount` command. See the `mount_nfs(1M)` man page. Under most conditions, allow the negotiation to select the best options.

File Transfer Size Negotiation

The file transfer size establishes the size of the buffers that are used when transferring data between the client and the server. In general, larger transfer sizes are better. The NFS version 3 protocol has an unlimited transfer size. However, starting with the Solaris 2.6 release, the software bids a default buffer size of 32 Kbytes. The client can bid a smaller transfer size at mount time if needed, but under most conditions this bid is not necessary.

The transfer size is not negotiated with systems that use the NFS version 2 protocol. Under this condition, the maximum transfer size is set to 8 Kbytes.

You can use the `-rsize` and `-wsize` options to set the transfer size manually with the `mount` command. You might need to reduce the transfer size for some PC clients. Also, you can increase the transfer size if the NFS server is configured to use larger transfer sizes.

How File Systems Are Mounted

When a client needs to mount a file system from a server, the client must obtain a file handle from the server. The file handle must correspond to the file system. This process requires that several transactions occur between the client and the server. In this example, the client is attempting to mount `/home/terry` from the server. A snoop trace for this transaction follows.

```
client -> server PORTMAP C GETPORT prog=100005 (MOUNT) vers=3 proto=UDP
server -> client PORTMAP R GETPORT port=33492
client -> server MOUNT3 C Null
server -> client MOUNT3 R Null
client -> server MOUNT3 C Mount /export/home9/terry
server -> client MOUNT3 R Mount OK FH=9000 Auth=unix
client -> server PORTMAP C GETPORT prog=100003 (NFS) vers=3 proto=TCP
server -> client PORTMAP R GETPORT port=2049
client -> server NFS C NULL3
server -> client NFS R NULL3
client -> server NFS C FSINFO3 FH=9000
server -> client NFS R FSINFO3 OK
client -> server NFS C GETATTR3 FH=9000
server -> client NFS R GETATTR3 OK
```

In this trace, the client first requests the mount port number from the portmap service on the NFS server. After the client receives the mount port number (33492), that number is used to ping the service on the server. After the client has determined that a service is running on that port number, the client then makes a mount request. When the server responds to this request, the server includes the file handle for the file system (9000) being mounted. The client then sends a request for the NFS port number. When the client receives the number from the server, the client pings the NFS service (nfsd). Also, the client requests NFS information about the file system that uses the file handle.

In the following trace, the client is mounting the file system with the `public` option.

```
client -> server NFS C LOOKUP3 FH=0000 /export/home9/terry
server -> client NFS R LOOKUP3 OK FH=9000
client -> server NFS C FSINFO3 FH=9000
server -> client NFS R FSINFO3 OK
client -> server NFS C GETATTR3 FH=9000
server -> client NFS R GETATTR3 OK
```

By using the default public file handle (which is 0000), all of the transactions to obtain information from the portmap service and to determine the NFS port number are skipped.

Effects of the `-public` Option and NFS URLs When Mounting

Using the `-public` option can create conditions that cause a mount to fail. Adding an NFS URL can also confuse the situation. The following list describes the specifics of how a file system is mounted when you use these options.

Public option with NFS URL – Forces the use of the public file handle. The mount fails if the public file handle is not supported.

Public option with regular path – Forces the use of the public file handle. The mount fails if the public file handle is not supported.

NFS URL only – Use the public file handle if this file handle is enabled on the NFS server. If the mount fails when using the public file handle, then try the mount with the MOUNT protocol.

Regular path only – Do not use the public file handle. The MOUNT protocol is used.

Client-Side Failover

By using client-side failover, an NFS client can switch to another server if the server that supports a replicated file system becomes unavailable. The file system can become unavailable under one of the following circumstances.

- If the server that the file system is connected to crashes
- If the server is overloaded
- If a network fault occurs

The failover, under these conditions, is normally transparent to the user. Thus, the failover can occur at any time without disrupting the processes that are running on the client.

Failover requires that the file system be mounted read-only. The file systems must be identical for the failover to occur successfully. See “What Is a Replicated File System?” on page 225 for a description of what makes a file system identical. A static file system or a file system that is not changed often is the best candidate for failover.

You cannot use file systems that are mounted by using CacheFS with failover. Extra information is stored for each CacheFS file system. This information cannot be updated during failover, so only one of these two features can be used when mounting a file system.

The number of replicas that need to be established for every file system depends on many factors. Ideally, you should have a minimum of two servers. Each server should support multiple subnets. This setup is better than having a unique server on each subnet. The process requires that each listed server be checked. Therefore, if more servers are listed, each mount is slower.

Failover Terminology

To fully comprehend the process, you need to understand two terms.

- *failover* – The process of selecting a server from a list of servers that support a replicated file system. Normally, the next server in the sorted list is used, unless it fails to respond.
- *remap* – To make use of a new server. Through normal use, the clients store the path name for each active file on the remote file system. During the remap, these path names are evaluated to locate the files on the new server.

What Is a Replicated File System?

For the purposes of failover, a file system can be called a *replica* when each file is the same size and has the same vnode type as the original file system. Permissions, creation dates, and other file attributes are not considered. If the file size or vnode types are different, the remap fails and the process hangs until the old server becomes available.

You can maintain a replicated file system by using *rdist*, *cpio*, or another file transfer mechanism. Because updating the replicated file systems causes inconsistency, follow these suggestions for best results:

- Rename the old version of the file before installing a new version of the file.

- Run the updates at night when client usage is low.
- Keep the updates small.
- Minimize the number of copies.

Failover and NFS Locking

Some software packages require read locks on files. To prevent these products from breaking, read locks on read-only file systems are allowed but are visible to the client side only. The locks persist through a remap because the server does not “know” about the locks. Because the files should not change, you do not need to lock the file on the server side.

Large Files

Starting with 2.6, the Solaris release supports files that are over 2 Gbytes. By default, UFS file systems are mounted with the `-largefiles` option to support the new capability. Previous releases cannot handle files of this size. See “How to Disable Large Files on an NFS Server” on page 159 for instructions.

If the server’s file system is mounted with the `-largefiles` option, a Solaris 2.6 NFS client can access large files without the need for changes. However, not all 2.6 commands can handle these large files. See `largefile(5)` for a list of the commands that can handle the large files. Clients that cannot support the NFS version 3 protocol with the large file extensions cannot access any large files. Although clients that run the Solaris 2.5 release can use the NFS version 3 protocol, large file support was not included in that release.

How NFS Server Logging Works

NFS server logging provides records of NFS reads and writes, as well as operations that modify the file system. This data can be used to track access to information. In addition, the records can provide a quantitative way to measure interest in the information.

When a file system with logging enabled is accessed, the kernel writes raw data into a buffer file. This data includes the following:

- A timestamp
- The client IP address
- The UID of the requester
- The file handle of the file or directory object that is being accessed
- The type of operation that occurred

The `nfslogd` daemon converts this raw data into ASCII records that are stored in log files. During the conversion, the IP addresses are modified to host names and the UIDs are modified to logins if the name service that is enabled can find matches. The file

handles are also converted into path names. To accomplish the conversion, the daemon tracks the file handles and stores information in a separate file handle-to-path table. That way, the path does not have to be re-identified each time a file handle is accessed. Because no changes to the mappings are made in the file handle-to-path table if `nfslogd` is turned off, you must keep the daemon running.

How the WebNFS Service Works

The WebNFS service makes files in a directory available to clients by using a public file handle. A file handle is an address that is generated by the kernel that identifies a file for NFS clients. The *public file handle* has a predefined value, so the server does not need to generate a file handle for the client. The ability to use this predefined file handle reduces network traffic by eliminating the MOUNT protocol. This ability should also accelerate processes for the clients.

By default, the public file handle on an NFS server is established on the root file system. This default provides WebNFS access to any clients that already have mount privileges on the server. You can change the public file handle to point to any file system by using the `share` command.

When the client has the file handle for the file system, a LOOKUP is run to determine the file handle for the file to be accessed. The NFS protocol allows the evaluation of only one path name component at a time. Each additional level of directory hierarchy requires another LOOKUP. A WebNFS server can evaluate an entire path name with a single multicomponent lookup transaction when the LOOKUP is relative to the public file handle. Multicomponent lookup enables the WebNFS server to deliver the file handle to the desired file without exchanging the file handles for each directory level in the path name.

In addition, an NFS client can initiate concurrent downloads over a single TCP connection. This connection provides quick access without the additional load on the server that is caused by setting up multiple connections. Although web browser applications support concurrent downloading of multiple files, each file has its own connection. By using one connection, the WebNFS software reduces the overhead on the server.

If the final component in the path name is a symbolic link to another file system, the client can access the file if the client already has access through normal NFS activities.

Normally, an NFS URL is evaluated relative to the public file handle. The evaluation can be changed to be relative to the server's root file system by adding an additional slash to the beginning of the path. In this example, these two NFS URLs are equivalent if the public file handle has been established on the `/export/ftp` file system.

```
nfs://server/junk  
nfs://server//export/ftp/junk
```

How WebNFS Security Negotiation Works

The Solaris 8 release includes a new protocol so a WebNFS client can negotiate a selected security mechanism with a WebNFS server. The new protocol uses security negotiation multicomponent lookup, which is an extension to the multicomponent lookup that was used in earlier versions of the WebNFS protocol.

The WebNFS client initiates the process by making a regular multicomponent lookup request by using the public file handle. Because the client has no knowledge of how the path is protected by the server, the default security mechanism is used. If the default security mechanism is not sufficient, the server replies with an `AUTH_TOOWEAK` error. This reply indicates that the default mechanism is not valid. The client needs to use a stronger default mechanism.

When the client receives the `AUTH_TOOWEAK` error, the client sends a request to the server to determine which security mechanisms are required. If the request succeeds, the server responds with an array of security mechanisms that are required for the specified path. Depending on the size of the array of security mechanisms, the client might have to make more requests to obtain the complete array. If the server does not support WebNFS security negotiation, the request fails.

After a successful request, the WebNFS client selects the first security mechanism from the array that the client supports. The client then issues a regular multicomponent lookup request by using the selected security mechanism to acquire the file handle. All subsequent NFS requests are made by using the selected security mechanism and the file handle.

WebNFS Limitations With Web Browser Use

Several functions that a web site that uses HTTP can provide are not supported by the WebNFS software. These differences stem from the fact that the NFS server only sends the file, so any special processing must be done on the client. If you need to have one web site configured for both WebNFS and HTTP access, consider the following issues:

- NFS browsing does not run CGI scripts. So, a file system with an active web site that uses many CGI scripts might not be appropriate for NFS browsing.
- The browser might start different viewers in order to handle files in different file formats. Accessing these files through an NFS URL starts an external viewer if the file type can be determined by the file name. The browser should recognize any file name extension for a standard MIME type when an NFS URL is used. As an explanation, the WebNFS software does not check inside the file to determine the file type. So, the only way to determine a file type is by the file name extension.
- NFS browsing cannot utilize server-side image maps (clickable images). However, NFS browsing can utilize client-side image maps (clickable images) because the URLs are defined with the location. No additional response is required from the document server.

Secure NFS System

The NFS environment is a powerful way and convenient way to share file systems on a network of different computer architectures and operating systems. However, the same features that make sharing file systems through NFS operation convenient also pose some security problems. Historically, most NFS implementations have used UNIX (or AUTH_SYS) authentication, but stronger authentication methods such as AUTH_DH have also been available. When using UNIX authentication, an NFS server authenticates a file request by authenticating the computer that makes the request, but not the user. Therefore, a client user can run `su` and impersonate the owner of a file. If DH authentication is used, the NFS server authenticates the user, making this sort of impersonation much harder.

With root access and knowledge of network programming, anyone can introduce arbitrary data into the network and extract any data from the network. The most dangerous attacks are those attacks that involve the introduction of data. An example is the impersonation of a user by generating the right packets or by recording “conversations” and replaying them later. These attacks affect data integrity. Attacks that involve passive eavesdropping—merely listening to network traffic without impersonating anybody—are not as dangerous, as data integrity is not compromised. Users can protect the privacy of sensitive information by encrypting data that is sent over the network.

A common approach to network security problems is to leave the solution to each application. A better approach is to implement a standard authentication system at a level that covers all applications.

The Solaris operating environment includes an authentication system at the level of remote procedure call (RPC)—the mechanism on which NFS operation is built. This system, known as Secure RPC, greatly improves the security of network environments and provides additional security to services such as the NFS system. When the NFS system uses the facilities that are provided by Secure RPC, it is known as a Secure NFS system.

Secure RPC

Secure RPC is fundamental to the Secure NFS system. The goal of Secure RPC is to build a system that is at minimum as secure as a time-sharing system. In a time-sharing system all users share a single computer. A time-sharing system authenticates a user through a login password. With data encryption standard (DES) authentication, the same authentication process is completed. Users can log in on any remote computer just as users can log in on a local terminal. The users’ login passwords are their passports to network security. In a time-sharing environment, the system administrator has an ethical obligation not to change a password to impersonate someone. In Secure RPC, the network administrator is trusted not to alter entries in a database that stores *public keys*.

You need to be familiar with two terms to understand an RPC authentication system: credentials and verifiers. Using ID badges as an example, the credential is what identifies a person: a name, address, birthday, and so on. The verifier is the photo that is attached to the badge. You can be sure the badge has not been stolen by checking the photo on the badge against the person who is carrying the badge. In RPC, the client process sends both a credential and a verifier to the server with each RPC request. The server sends back only a verifier because the client already “knows” the server’s credentials.

RPC’s authentication is open ended, which means that a variety of authentication systems can be plugged into it, such as UNIX, DH, and KERB.

When UNIX authentication is used by a network service, the credentials contain the client’s host name, UID, GID, and group-access list. However, the verifier contains nothing. Because no verifier exists, a superuser could falsify appropriate credentials by using commands such as `su`. Another problem with UNIX authentication is that UNIX authentication assumes all computers on a network are UNIX computers. UNIX authentication breaks down when applied to other operating systems in a heterogeneous network.

To overcome the problems of UNIX authentication, Secure RPC uses DH authentication.

DH Authentication

DH authentication uses the Data Encryption Standard (DES) and Diffie-Hellman public-key cryptography to authenticate both users and computers in the network. DES is a standard encryption mechanism. Diffie-Hellman public-key cryptography is a cipher system that involves two keys: one public and one secret. The public keys and secret keys are stored in the namespace. NIS stores the keys in the public-key map. These maps contain the public key and secret key for all potential users. See the *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)* for more information on how to set up the maps.

The security of DH authentication is based on a sender’s ability to encrypt the current time, which the receiver can then decrypt and check against its own clock. The timestamp is encrypted with DES. The requirements for this scheme to work are as follows:

- The two agents must agree on the current time.
- The sender and receiver must be using the same encryption key.

If a network runs a time-synchronization program, the time on the client and the server is synchronized automatically. If a time-synchronization program is not available, timestamps can be computed by using the server’s time instead of the network time. The client asks the server for the time before starting the RPC session, then computes the time difference between its own clock and the server’s. This

difference is used to offset the client's clock when computing timestamps. If the client and server clocks get out of synchronization to the point where the server begins to reject the client's requests, the DH authentication system on the client resynchronizes with the server.

The client and server arrive at the same encryption key by generating a random *conversation key*, also known as the *session key*, and by using public-key cryptography to deduce a *common key*. The common key is a key that only the client and server are capable of deducing. The conversation key is used to encrypt and decrypt the client's timestamp. The common key is used to encrypt and decrypt the conversation key.

KERB Authentication

Kerberos is an authentication system that was developed at MIT. Encryption in Kerberos is based on DES. Kerberos support is no longer supplied as part of Secure RPC, but a server-side and client-side implementation is included with the Solaris 9 release. See "Introduction to SEAM" in *System Administration Guide: Security Services* for more information about the Solaris 9 implementation of Kerberos Authentication.

Using Secure RPC With NFS

Be aware of the following points if you plan to use Secure RPC:

- If a server crashes when no one is around (after a power failure, for example), all the secret keys that are stored on the system are deleted. Now no process can access secure network services or mount an NFS file system. The important processes during a reboot are usually run as `root`. Therefore, these processes would work if `root`'s secret key were stored away, but nobody is available to type the password that decrypts it. `keylogin -r` allows `root` to store the clear secret key in `/etc/.rootkey`, which `keyserv` reads.
- Some systems boot in single-user mode, with a `root` login shell on the console and no password prompt. Physical security is imperative in such cases.
- Diskless computer booting is not totally secure. Somebody could impersonate the boot server and boot a devious kernel that, for example, makes a record of your secret key on a remote computer. The Secure NFS system provides protection only after the kernel and the key server are running. Otherwise, no way exists to authenticate the replies that are given by the boot server. This limitation could be a serious problem, but the limitation requires a sophisticated attack, using kernel source code. Also, the crime would leave evidence. If you polled the network for boot servers, you would discover the devious boot server's location.
- Most `setuid` programs are owned by `root`. If the secret key for `root` is stored in `/etc/.rootkey`, these programs behave as they always have. If a `setuid` program is owned by a user, however, the `setuid` program might not always work. For example, suppose that a `setuid` program is owned by `dave` and `dave` has not logged into the computer since it booted. The program would not be able to access

secure network services.

- If you log in to a remote computer (using `login`, `rlogin`, or `telnet`) and use `keylogin` to gain access, you give access to your account. The reason is that your secret key is passed to that computer's key server, which then stores your secret key. This process is only a concern if you do not trust the remote computer. If you have doubts, however, do not log in to a remote computer if the remote computer requires a password. Instead, use the NFS environment to mount file systems that are shared by the remote computer. As an alternative, you can use `keylogout` to delete the secret key from the key server.
- If a home directory is shared with the `-o sec=dh` option, remote logins can be a problem. If the `/etc/hosts.equiv` or `~/.rhosts` files are not set to prompt for a password, the login succeeds. However, the users cannot access their home directories because no authentication has occurred locally. If the user is prompted for a password, the user has access to his or her home directory if the password matches the network password.

Autofs Maps

Autofs uses three types of maps:

- Master map
- Direct maps
- Indirect maps

Master Autofs Map

The `auto_master` map associates a directory with a map. The map is a master list that specifies all the maps that autofs should check. The following example shows what an `auto_master` file could contain.

EXAMPLE 16-1 Sample `/etc/auto_master` File

```
# Master map for automounter
#
+auto_master
/net          -hosts          -nosuid,nobrowse
/home         auto_home       -nobrowse
/xfn          -xfn
/-            auto_direct    -ro
```

This example shows the generic `auto_master` file with one addition for the `auto_direct` map. Each line in the master map `/etc/auto_master` has the following syntax:

mount-point map-name [mount-options]

<i>mount-point</i>	<i>mount-point</i> is the full (absolute) path name of a directory. If the directory does not exist, autofs creates the directory if possible. If the directory exists and is not empty, mounting on the directory hides its contents. In this situation, autofs issues a warning. The notation / - as a mount point indicates that the map in question is a direct map, and no particular mount point is associated with the map as a whole.
<i>map-name</i>	<i>map-name</i> is the map autofs uses to find directions to locations, or mount information. If the name is preceded by a slash (/), autofs interprets the name as a local file. Otherwise, autofs searches for the mount information by using the search that is specified in the name-service switch configuration file (/etc/nsswitch.conf). Special maps are also used for /net and /xfn. See "Mount Point /net" on page 234 and "Mount Point /xfn" on page 234 for more information.
<i>mount-options</i>	<i>mount-options</i> is an optional, comma-separated list of options that apply to the mounting of the entries that are specified in map-name, unless the entries in map-name list other options. Options for each specific type of file system are listed in the mount man page for that file system. For example, see the mount_nfs(1M) man page for NFS-specific mount options. For NFS-specific mount points, the bg (background) and fg (foreground) options do not apply.

A line that begins with # is a comment. All the text that follows until the end of the line is ignored.

To split long lines into shorter ones, put a backslash (\) at the end of the line. The maximum number of characters of an entry is 1024.

Note – If the same mount point is used in two entries, the first entry is used by the automount command. The second entry is ignored.

Mount Point /home

The mount point /home is the directory under which the entries that are listed in /etc/auto_home (an indirect map) are to be mounted.

Note – Autofs runs on all computers and supports /net and /home (automounted home directories) by default. These defaults can be overridden by entries in the NIS auto.master map or NIS+ auto_master table, or by local editing of the /etc/auto_master file.

Mount Point /net

Autofs mounts under the directory /net all the entries in the special map -hosts. The map is a built-in map that uses only the hosts database. Suppose that the computer gumbo is in the hosts database and it exports any of its file systems. The following command changes the current directory to the root directory of the computer gumbo.

```
% cd /net/gumbo
```

Autofs can mount only the *exported* file systems of host gumbo, that is, those on a server available to network users as opposed to those on a local disk. Therefore, all the files and directories on gumbo might not be available through /net/gumbo.

With the /net method of access, the server name is in the path and is location dependent. If you want to move an exported file system from one server to another, the path might no longer work. Instead, you should set up an entry in a map specifically for the file system you want rather than use /net.

Note – Autofs checks the server’s export list only at mount time. After a server’s file systems are mounted, autofs does not check with the server again until the server’s file systems are automatically unmounted. Therefore, newly exported file systems are not “seen” until the file systems on the client are unmounted and then remounted.

Mount Point /xfsn

This mount point provides the autofs directory structure for the resources that are shared through the FNS namespace. See the *System Administration Guide: Naming and Directory Services (FNS and NIS+)* for more information about FNS.

Direct Autofs Maps

A direct map is an automount point. With a direct map, a direct association exists between a mount point on the client and a directory on the server. Direct maps have a full path name and indicate the relationship explicitly. The following is a typical /etc/auto_direct map:

```
/usr/local      -ro \
  /bin           ivy:/export/local/sun4 \
  /share         ivy:/export/local/share \
  /src           ivy:/export/local/src
/usr/man        -ro oak:/usr/man \
                rose:/usr/man \
                willow:/usr/man
/usr/games      -ro peach:/usr/games
/usr/spool/news -ro pine:/usr/spool/news \
```

willow:/var/spool/news

Lines in direct maps have the following syntax:

key [*mount-options*] *location*

<i>key</i>	<i>key</i> is the path name of the mount point in a direct map.
<i>mount-options</i>	<i>mount-options</i> is the options that you want to apply to this particular mount. These options are required only if the options differ from the map default. Options for each specific type of file system are listed in the mount man page for that file system. For example, see the mount_cacheefs(1M) man page for CacheFS specific mount options.
<i>location</i>	<i>location</i> is the location of the file system. One or more file systems are specified as <i>server:pathname</i> for NFS file systems or <i>:devicename</i> for High Sierra file systems (HSFS). Note – The <i>pathname</i> should not include an automounted mount point. The <i>pathname</i> should be the actual absolute path to the file system. For instance, the location of a home directory should be listed as <i>server:/export/home/username</i> , not as <i>server:/home/username</i> .

As in the master map, a line that begins with # is a comment. All the text that follows until the end of the line is ignored. Put a backslash at the end of the line to split long lines into shorter ones.

Of all the maps, the entries in a direct map most closely resemble the corresponding entries in /etc/vfstab. An entry might appear in /etc/vfstab as follows:

```
dancer:/usr/local - /usr/local/tmp nfs - yes ro
```

The equivalent entry appears in a direct map as follows:

```
/usr/local/tmp      -ro      dancer:/usr/local
```

Note – No concatenation of options occurs between the automounter maps. Any options that are added to an automounter map override all options that are listed in maps that are searched earlier. For instance, options that are included in the auto_master map would be overridden by corresponding entries in any other map.

See “How Autofs Selects the Nearest Read-Only Files for Clients (Multiple Locations)” on page 242 for other important features that are associated with this type of map.

Mount Point /-

In Example 16-1, the mount point /- tells autofs not to associate the entries in `auto_direct` with any specific mount point. Indirect maps use mount points that are defined in the `auto_master` file. Direct maps use mount points that are specified in the named map. Remember, in a direct map the key, or mount point, is a full path name.

An NIS or NIS+ `auto_master` file can have only one direct map entry because the mount point must be a unique value in the namespace. An `auto_master` file that is a local file can have any number of direct map entries if entries are not duplicated.

Indirect Autofs Maps

An indirect map uses a substitution value of a key to establish the association between a mount point on the client and a directory on the server. Indirect maps are useful for accessing specific file systems, such as home directories. The `auto_home` map is an example of an indirect map.

Lines in indirect maps have the following general syntax:

key [*mount-options*] *location*

<i>key</i>	<i>key</i> is a simple name without slashes in an indirect map.
<i>mount-options</i>	<i>mount-options</i> is the options that you want to apply to this particular mount. These options are required only if the options differ from the map default. Options for each specific type of file system are listed in the mount man page for that file system. For example, see the <code>mount_nfs(1M)</code> man page for NFS-specific mount options.
<i>location</i>	<i>location</i> is the location of the file system. One or more file systems are specified as <i>server:pathname</i> . Note – The <i>pathname</i> should not include an automounted mount point. The <i>pathname</i> should be the actual absolute path to the file system. For instance, the location of a directory should be listed as <i>server:/usr/local</i> , not as <i>server:/net/server/usr/local</i> .

As in the master map, a line that begins with # is a comment. All the text that follows until the end of the line is ignored. Put a backslash (\) at the end of the line to split long lines into shorter ones. Example 16-1 shows an `auto_master` map that contains the following entry:

```
/home      auto_home      -nobrowse
```

`auto_home` is the name of the indirect map that contains the entries to be mounted under `/home`. A typical `auto_home` map might contain the following:

Composed September 5, 2002

```
david          willow:/export/home/david
rob            cypress:/export/home/rob
gordon         poplar:/export/home/gordon
rajan          pine:/export/home/rajan
tammy          apple:/export/home/tammy
jim            ivy:/export/home/jim
linda -rw,nosuid peach:/export/home/linda
```

As an example, assume that the previous map is on host oak. Suppose that the user linda has an entry in the password database that specifies her home directory as /home/linda. Whenever linda logs in to computer oak, autofs mounts the directory /export/home/linda that resides on the computer peach. Her home directory is mounted read-write, nosuid.

Assume the following conditions occur: User linda's home directory is listed in the password database as /home/linda. Anybody, including Linda, has access to this path from any computer that is set up with the master map referring to the map in the previous example.

Under these conditions, user linda can run `login` or `rlogin` on any of these computers and have her home directory mounted in place for her.

Furthermore, now Linda can also type the following command:

```
% cd ~david
```

autofs mounts David's home directory for her (if all permissions allow).

Note – No concatenation of options occurs between the automounter maps. Any options that are added to an automounter map override all options that are listed in maps that are searched earlier. For instance, options that are included in the `auto_master` map are overridden by corresponding entries in any other map.

On a network without a name service, you have to change all the relevant files (such as `/etc/passwd`) on all systems on the network to allow Linda access to her files. With NIS, make the changes on the NIS master server and propagate the relevant databases to the slave servers. On a network that is running NIS+, propagating the relevant databases to the slave servers is done automatically after the changes are made.

How Autofs Works

Autofs is a client-side service that automatically mounts the appropriate file system. When a client attempts to access a file system that is not presently mounted, the autofs file system intercepts the request and calls `automountd` to mount the requested directory. The `automountd` daemon locates the directory, mounts it within autofs, and replies. On receiving the reply, autofs allows the waiting request to proceed. Subsequent references to the mount are redirected by the autofs. No further participation is required by `automountd` until the file system is automatically unmounted by autofs after a period of inactivity.

The components that work together to accomplish automatic mounting are the following:

- The `automount` command
- The autofs file system
- The `automountd` daemon

The `automount` command, which is called at system startup time, reads the master map file `auto_master` to create the initial set of autofs mounts. These autofs mounts are not automatically mounted at startup time. These mounts are points under which file systems are mounted in the future. These points are also known as trigger nodes.

After the autofs mounts are set up, these mounts can trigger file systems to be mounted under them. For example, when autofs receives a request to access a file system that is not currently mounted, autofs calls `automountd`, which actually mounts the requested file system.

After initially mounting autofs mounts, the `automount` command is used to update autofs mounts as necessary. The command compares the list of mounts in the `auto_master` map with the list of mounted file systems in the mount table file `/etc/mnttab` (formerly `/etc/mtab`). `automount` then makes the appropriate changes. This process allows system administrators to change mount information within `auto_master` and have those changes used by the autofs processes without having to stop and restart the autofs daemon. After the file system is mounted, further access does not require any action from `automountd` until the file system is automatically unmounted.

Unlike `mount`, `automount` does not read the `/etc/vfstab` file (which is specific to each computer) for a list of file systems to mount. The `automount` command is controlled within a domain and on computers through the namespace or local files.

The following is a simplified overview of how autofs works.

The automount daemon `automountd` starts at boot time from the `/etc/init.d/autofs` script (see Figure 16–1). This script also runs the `automount` command, which reads the master map and installs autofs mount points. See “How Autofs Starts the Navigation Process (Master Map)” on page 240 for more information.




FIGURE 16–1 `/etc/init.d/autofs` Script Starts automount

Autofs is a kernel file system that supports automatic mounting and unmounting.

When a request is made to access a file system at an autofs mount point, the following occurs:

1. Autofs intercepts the request.
2. Autofs sends a message to the `automountd` for the requested file system to be mounted.
3. `automountd` locates the file system information in a map, creates the trigger nodes, and performs the mount.
4. Autofs allows the intercepted request to proceed.
5. Autofs unmounts the file system after a period of inactivity.

Note – Mounts that are managed through the autofs service should not be manually mounted or unmounted. Even if the operation is successful, the autofs service does not check that the object has been unmounted, resulting in possible inconsistencies. A reboot clears all of the autofs mount points.

How Autofs Navigates Through the Network (Maps)

Autofs searches a series of maps to navigate through the network. Maps are files that contain information such as the password entries of all users on a network or the names of all host computers on a network. Effectively, the maps contain network-wide equivalents of UNIX administration files. Maps are available locally or through a

network name service such as NIS or NIS+. You create maps to meet the needs of your environment by using the Solstice System Management Tools. See “Modifying How Autofs Navigates the Network (Modifying Maps)” on page 247.

How Autofs Starts the Navigation Process (Master Map)

The automount command reads the master map at system startup. Each entry in the master map is a direct map name or an indirect map name, its path, and its mount options, as shown in Figure 16–2. The specific order of the entries is not important. automount compares entries in the master map with entries in the mount table to generate a current list.




FIGURE 16–2 Navigation Through the Master Map

Autofs Mount Process

What the autofs service does when a mount request is triggered depends on how the automounter maps are configured. The mount process is generally the same for all mounts. However, the final result changes with the mount point that is specified and the complexity of the maps. Starting with the Solaris 2.6 release, the mount process has also been changed to include the creation of the trigger nodes.

Simple Autofs Mount

To help explain the autofs mount process, assume that the following files are installed.

```
$ cat /etc/auto_master
# Master map for automounter
#
+auto_master
/net          -hosts          -nosuid,nobrowse
/home         auto_home       -nobrowse
/xfn          -xfn
```


Composed September 5, 2002

```
/share      auto_share
$ cat /etc/auto_share
# share directory map for automounter
#
ws          gumbo:/export/share/ws
```

When the `/share` directory is accessed, the `autofs` service creates a trigger node for `/share/ws`, which can be seen in `/etc/mnttab` as an entry that resembles the following entry:

```
-hosts /share/ws      autofs  nosuid,nobrowse,ignore,nest,dev=###
```

When the `/share/ws` directory is accessed, the `autofs` service completes the process with these steps:

1. Pings the server's mount service to see if the service is alive.
2. Mounts the requested file system under `/share`. Now, the `/etc/mnttab` file contains the following entries:

```
-hosts /share/ws      autofs  nosuid,nobrowse,ignore,nest,dev=###
gumbo:/export/share/ws /share/ws  nfs    nosuid,dev=####  #####
```

Hierarchical Mounting

When multiple layers are defined in the automounter files, the mount process becomes more complex. Suppose that you expand the `/etc/auto_shared` file from the previous example to contain the following:

```
# share directory map for automounter
#
ws          /          gumbo:/export/share/ws
           /usr       gumbo:/export/share/ws/usr
```

The mount process is basically the same as the previous example when the `/share/ws` mount point is accessed. In addition, a trigger node to the next level (`/usr`) is created in the `/share/ws` file system so that the next level can be mounted if it is accessed. In this example, `/export/share/ws/usr` must exist on the NFS server in order for the trigger node to be created.



Caution – Do not use the `-soft` option when specifying hierarchical layers. Refer to “Autofs Unmounting” on page 242 for an explanation of this limitation.

Autofs Unmounting

The unmounting that occurs after a certain amount of idle time is from the bottom up (reverse order of mounting). If one of the directories at a higher level in the hierarchy is busy, only file systems below that directory are unmounted. During the unmounting process, any trigger nodes are removed and then the file system is unmounted. If the file system is busy, the unmount fails and the trigger nodes are reinstalled.



Caution – Do not use the `-soft` option when specifying hierarchical layers. If the `-soft` option is used, requests to reinstall the trigger nodes can time out. The failure to reinstall the trigger nodes leaves no access to the next level of mounts. The only way to clear this problem is to have the automounter unmount all of the components in the hierarchy. The automounter can complete the unmount either by waiting for the file systems to be automatically unmounted or by rebooting the system.

How Autofs Selects the Nearest Read-Only Files for Clients (Multiple Locations)

The example direct map contains the following:

```
/usr/local      -ro \
  /bin          ivy:/export/local/sun4\
  /share        ivy:/export/local/share\
  /src          ivy:/export/local/src
/usr/man        -ro oak:/usr/man \
               rose:/usr/man \
               willow:/usr/man
/usr/games      -ro peach:/usr/games
/usr/spool/news -ro pine:/usr/spool/news \
               willow:/var/spool/news
```

The mount points `/usr/man` and `/usr/spool/news` list more than one location, three locations for the first mount point, two locations for the second mount point. This means any of the replicated locations can provide the same service to any user. This procedure is sensible only when you mount a file system that is read-only, as you must have some control over the locations of files that you write or modify. You don't want to modify files on one server on one occasion and, minutes later, modify the "same" file on another server. The benefit is that the best available server is used automatically without any effort that is required by the user.

If the file systems are configured as replicas (see "What Is a Replicated File System?" on page 225), the clients have the advantage of using failover. Not only is the best server automatically determined, but if that server becomes unavailable, the client automatically uses the next-best server. Failover was first implemented in the Solaris 2.6 release.

An example of a good file system to configure as a replica is man pages. In a large network, more than one server can export the current set of man pages. Which server you mount the man pages from does not matter if the server is running and exporting its file systems. In the previous example, multiple mount locations are expressed as a list of mount locations in the map entry.

```
/usr/man -ro oak:/usr/man rose:/usr/man willow:/usr/man
```

In this example, you can mount the man pages from the servers *oak*, *rose*, or *willow*. Which server is best depends on a number of factors, including the following:

- The number of servers that support a particular NFS protocol level
- The proximity of the server
- The weighting

During the sorting process, a count is taken of the number of servers that support the NFS version 2 and NFS version 3 protocols. Whichever protocol is supported on the most servers becomes the protocol that is supported by default. This selection provides the client with the maximum number of servers to depend on.

After the largest subset of servers with the same protocol version is found, that server list is sorted by proximity. Servers on the local subnet are given preference over servers on a remote subnet. The closest server is given preference, which reduces latency and network traffic. Figure 16–3 illustrates server proximity.




FIGURE 16–3 Server Proximity

If several servers that support the same protocol are on the local subnet, the time to connect to each server is determined and the fastest server is used. The sorting can also be influenced by using weighting (see “Autofs and Weighting” on page 244).

If version 3 servers are more abundant, the sorting process becomes more complex. Normally, servers on the local subnet are given preference over servers on a remote subnet. A version 2 server can complicate matters, as a version 2 server might be closer than the nearest version 3 server. If a version 2 server is on the local subnet and the closest version 3 server is on a remote subnet, the version 2 server is given preference. This preference is only checked if more version 3 servers exist than version 2 servers. If more version 2 servers exist, only a version 2 server is selected.

With failover, the sorting is checked once at mount time to select one server from which to mount, and again anytime the mounted server becomes unavailable. Multiple locations are useful in an environment where individual servers might not export their file systems temporarily.

This feature is particularly useful in a large network with many subnets. Autofs chooses the nearest server and therefore confines NFS network traffic to a local network segment. In servers with multiple network interfaces, list the host name that is associated with each network interface as if the interface were a separate server. Autofs selects the nearest interface to the client.

Autofs and Weighting

You can influence the selection of servers at the same proximity level by adding a weighting value to the autofs map. For example:

```
/usr/man -ro oak,rose(1),willow(2):/usr/man
```

The numbers in parentheses indicate a weighting. Servers without a weighting have a value of zero and, therefore, are most likely to be selected. The higher the weighting value, the lower the chance that the server is selected.

Note – All other server selection factors are more important than weighting. Weighting is only considered when selecting between servers with the same network proximity.

Variables in a Map Entry

You can create a client-specific variable by prefixing a dollar sign (\$) to its name. The variable helps you to accommodate different architecture types that are accessing the same file-system location. You can also use curly braces to delimit the name of the variable from appended letters or digits. Table 16–3 shows the predefined map variables.

TABLE 16-3 Predefined Map Variables

Variable	Meaning	Derived From	Example
ARCH	Architecture type	uname -m	sun4u
CPU	Processor type	uname -p	sparc
HOST	Host name	uname -n	dinky
OSNAME	Operating system name	uname -s	SunOS
OSREL	Operating system release	uname -r	5.8
OSVERS	Operating system version (version of the release)	uname -v	GENERIC

You can use variables anywhere in an entry line except as a key. For instance, suppose that you have a file server that exports binaries for SPARC and IA architectures from `/usr/local/bin/sparc` and `/usr/local/bin/x86` respectively. The clients can mount through a map entry such as the following:

```
/usr/local/bin      -ro      server:/usr/local/bin/$CPU
```

Now the same entry for all clients applies to all architectures.

Note – Most applications that are written for any of the sun4 architectures can run on all sun4 platforms, so the `-ARCH` variable is hard-coded to `sun4` instead of `sun4m`.

Maps That Refer to Other Maps

A map entry *+mapname* that is used in a file map causes automount to read the specified map as if it were included in the current file. If *mapname* is not preceded by a slash, autofs treats the map name as a string of characters and uses the name-service switch policy to find the map name. If the path name is an absolute path name, automount checks a local map of that name. If the map name starts with a dash (-), automount consults the appropriate built-in map, such as `xfn` or `hosts`.

This name service switch file contains an entry for autofs that is labeled as `automount`, which contains the order in which the name services are searched. The following file is an example of a name service switch file.

```
#
# /etc/nsswitch.nis:
#
# An example file that could be copied over to /etc/nsswitch.conf;
# it uses NIS (YP) in conjunction with files.
#
# "hosts:" and "services:" in this file are used only if the /etc/netconfig
```

Composed September 5, 2002

```
# file contains "switch.so" as a nametoaddr library for "inet" transports.
# the following two lines obviate the "+" entry in /etc/passwd and /etc/group.
passwd:      files nis
group:       files nis

# consult /etc "files" only if nis is down.
hosts:       nis [NOTFOUND=return] files
networks:    nis [NOTFOUND=return] files
protocols:   nis [NOTFOUND=return] files
rpc:         nis [NOTFOUND=return] files
ethers:      nis [NOTFOUND=return] files
netmasks:    nis [NOTFOUND=return] files
bootparams:  nis [NOTFOUND=return] files
publickey:   nis [NOTFOUND=return] files
netgroup:    nis
automount:   files nis
aliases:     files nis
# for efficient getservbyname() avoid nis
services:    files nis
```

In this example, the local maps are searched before the NIS maps. Therefore, you can have a few entries in your local `/etc/auto_home` map for the most commonly accessed home directories. You can then use the switch to fall back to the NIS map for other entries.

```
bill          cs.csc.edu:/export/home/bill
bonny         cs.csc.edu:/export/home/bonny
```

After consulting the included map, if no match is found, automount continues scanning the current map. Therefore, you can add more entries after a `+` entry.

```
bill          cs.csc.edu:/export/home/bill
bonny         cs.csc.edu:/export/home/bonny
+auto_home
```

The map that is included can be a local file or a built-in map. Remember, only local files can contain `+` entries.

```
+auto_home_finance    # NIS+ map
+auto_home_sales      # NIS+ map
+auto_home_engineering # NIS+ map
+ /etc/auto_mystuff    # local map
+auto_home             # NIS+ map
+-hosts                # built-in hosts map
```

Note – You cannot use `+` entries in NIS+ or NIS maps.

Executable Autofs Maps

You can create an autofs map that executes some commands to generate the autofs mount points. You could benefit from using an executable autofs map if you need to be able to create the autofs structure from a database or a flat file. The disadvantage to using an executable map is that the map needs to be installed on each host. An executable map cannot be included in either the NIS or the NIS+ name service.

The executable map must have an entry in the `auto_master` file.

```
/execute      auto_execute
```

Here is an example of an executable map:

```
#!/bin/ksh
#
# executable map for autofs
#

case $1 in
    src)  echo '--nosuid,hard bee:/export1' ;;
esac
```

For this example to work, the file must be installed as `/etc/auto_execute` and must have the executable bit set. Set permissions to 744. Under these circumstances, running the following command causes the `/export1` file system from `bee` to be mounted:

```
% ls /execute/src
```

Modifying How Autofs Navigates the Network (Modifying Maps)

You can modify, delete, or add entries to maps to meet the needs of your environment. As applications and other file systems that users require change their location, the maps must reflect those changes. You can modify autofs maps at any time. Whether your modifications are effective the next time `automountd` mounts a file system depends on which map you modify and what kind of modification you make.

Default Autofs Behavior With Name Services

Booting invokes autofs by using the `/etc/init.d/autofs` script and checks for the master `auto_master` map. Autofs is subject to the rules that are discussed subsequently.

Autofs uses the name service that is specified in the `automount` entry of the `/etc/nsswitch.conf` file. If NIS+ is specified, as opposed to local files or NIS, all map names are used as is. If NIS is selected and autofs cannot find a map that autofs

needs, but finds a map name that contains one or more underscores, the underscores are changed to dots. This change allows the old NIS file names to work. Then autofs checks the map again, as shown in Figure 16–4.




FIGURE 16–4 How Autofs Uses the Name Service

The screen activity for this session would resemble the following example.

```

$ grep /home /etc/auto_master
/home          auto_home

$ ypmatch brent auto_home
Can't match key brent in map auto_home. Reason: no such map in
server's domain.

$ ypmatch brent auto.home
diskus:/export/home/diskus1/&
  
```


If “files” is selected as the name service, all maps are assumed to be local files in the /etc directory. Autofs interprets a map name that begins with a slash (/) as local regardless of which name service autofs uses.

Autofs Reference

The remaining sections of this chapter describe more advanced autofs features and topics.

Metacharacters

Autofs recognizes some characters as having a special meaning. Some characters are used for substitutions, and some characters are used to protect other characters from the autofs map parser.

Ampersand (&)

If you have a map with many subdirectories specified, as in the following, consider using string substitutions.

```
john      willow:/home/john
mary      willow:/home/mary
joe       willow:/home/joe
able      pine:/export/able
baker     peach:/export/baker
```

You can use the ampersand character (&) to substitute the key wherever the key appears. If you use the ampersand, the previous map changes to the following:

```
john      willow:/home/&
mary      willow:/home/&
joe       willow:/home/&
able      pine:/export/&
baker     peach:/export/&
```

You could also use key substitutions in a direct map, in situations such as the following:

```
/usr/man                                willow,cedar,poplar:/usr/man
```

You can also simplify the entry further as follows:

```
/usr/man                                willow,cedar,poplar:&
```

Notice that the ampersand substitution uses the whole key string. Therefore, if the key in a direct map starts with a / (as it should), the slash is carried over, and you could not do, for example, the following:

```
/progs                &1, &2, &3:/export/src/progs
```

The reason is that autofs would interpret the example as the following:

```
/progs                /progs1, /progs2, /progs3:/export/src/progs
```

Asterisk ()*

You can use the universal substitute character, the asterisk (*), to match any key. You could mount the /export file system from all hosts through this map entry.

```
*                    &:/export
```

Each ampersand is substituted by the value of any given key. Autofs interprets the asterisk as an end-of-file character.

Special Characters

If you have a map entry that contains special characters, you might have to mount directories that have names that confuse the autofs map parser. The autofs parser is sensitive to names that contain colons, commas, spaces, and so on. These names should be enclosed in double quotations, as in the following:

```
/vms      -ro      vmsserver: - - - "rc0:dk1 - "  
/mac      -ro      gator:/ - "Mr Disk - "
```

CHAPTER **17**

SLP Topics

The section contains the following chapters on configuring and deploying Service Location Protocol (SLP) in the Solaris 9 operating environment.

Chapter 18	Information on SLP architecture and implementation
Chapter 19	SLP configuration considerations and the process to enable SLP
Chapter 20	Information and procedures for configuring SLP agents and processes
Chapter 21	Information on advertising SLP legacy services
Chapter 22	Tables that list SLP status codes and message types

Composed September 5, 2002

SLP (Overview)

The Service Location Protocol (SLP) provides a portable, platform-independent framework for the discovery and provisioning of SLP-enabled network services. This chapter describes the SLP architecture and the Solaris 9 implementation of SLP for IP intranets.

- “SLP Architecture” on page 253
- “SLP Implementation” on page 256

SLP Architecture

This section outlines the fundamental operation of SLP and describes agents and processes that are used in SLP administration.

SLP provides all of the following services automatically, with little or no configuration.

- Client application requests for information that is required to access a service
- Advertisement of services on network hardware devices or software servers; for example, printers, file servers, video cameras, and HTTP servers
- Managed recovery from primary server failures

In addition, you can do the following to administer and tune SLP operation if necessary.

- Organize services and users into *scopes* that are composed of logical or functional groups
- Enable SLP logging to monitor and troubleshoot the SLP operation on your network
- Adjust SLP timing parameters to enhance performance and scalability

- Configure SLP not to send and not to process multicast messages when SLP is deployed on networks that lack support for multicast routing
- Deploy SLP Directory Agents to improve scalability and performance

Summary of the SLP Design

SLP libraries inform network-aware agents that advertise services in order for those services to be discovered over a network. SLP agents maintain up-to-date information on the type and location of services. These agents can also use proxy registrations to advertise services that are not directly SLP enabled. For more information, see Chapter 21.

Client applications rely on SLP libraries that make requests directly to the agents that advertise services.

SLP Agents and Processes

The following table describes the SLP agents. For expanded definitions of these terms and other terms that are used in this volume, refer to the Glossary.

TABLE 18-1 SLP Agents

SLP Agent	Description
Directory Agent (DA)	Process that caches SLP advertisements that are registered by Service Agents (SAs). The DA forwards service advertisements to User Agents (UAs) on demand.
Service Agent (SA)	SLP agent that acts on behalf of a service to distribute service advertisements and to register the service with Directory Agents (DAs).
User Agent (UA)	SLP agent that acts on behalf of a user or application to obtain service advertisement information.
scope	An administrative or logical grouping of services.

The following figure shows the basic agents and processes that implement the SLP architecture. The figure represents a default deployment of SLP. No special configuration has been done. Only two agents are required: the UA and SA. The SLP framework allows the UA to multicast requests for services to the SA. The SA unicasts a reply to the UA. For example, when the UA sends a service request message, the SA responds with a service reply message. The service reply contains the location of services that match the client's requirements. Other requests and replies are possible for attributes and service types. For more information, see Chapter 22.




FIGURE 18-1 SLP Basic Agents and Processes

The following figure shows the basic agents and processes that implement the SLP architecture when a DA is deployed in the framework.




FIGURE 18-2 SLP Architectural Agents and Processes Implemented With a DA

When you deploy DAs, fewer messages are sent in the network and UAs can retrieve information much faster. DAs are essential when the size of a network increases or for situations in which there is no support for multicast routing. The DA serves as a cache for registered service advertisements. SAs send register messages (`SrvReg`) that list all the services they advertise to DAs. SAs then receive acknowledgments (`SrvAck`) in reply. The service advertisements are refreshed with the DA, or they expire according to the lifetime that is set for the advertisement. After a UA discovers a DA, the UA unicasts a request to the DA rather than multicasting requests to SAs.

For more information about Solaris SLP messages, refer to Chapter 22.

SLP Implementation

In the Solaris SLP implementation, the SLP SAs, UAs, DAs, SA servers, scopes, and other architectural components in Table 18–1 are partially mapped into `slpd` and partially into application processes. The SLP daemon, `slpd`, organizes certain off-host SLP interactions to do the following:

- Employ passive and active directory agent discovery in order to discover all DAs on the network
- Maintain an updated table of DAs for the use of the UAs and SAs on the local host
- Act as a proxy SA server for legacy service advertisements (proxy registration)

You can set the `net.slpisDA` property to also configure `slpd` to act as a DA. See Chapter 20.

For more information about the SLP daemon, see `slpd(1M)`.

In addition to `slpd`, the C/C++ and Java client libraries (`libslp.so` and `slp.jar`) enable access to the SLP framework for UA and SA clients. The client libraries provide the following features:

- Software that offers network services which can register and deregister service advertisements
- Client software that can request services by issuing queries for service advertisements
- The list of SLP scopes available for registration and requests

No special configuration is necessary to enable the inter-process communication between `slpd` and the client libraries that provide the previous services. You must, however, run the `slpd` process first before you load the client libraries in order for the libraries to function.

In the following figure, the SLP client library in the Service Provider Program employs SA functionality. The Service Provider Program uses the SLP client library to register and deregister services with `slpd`. The SLP client library in the Service Client Program employs UA functionality. The Service Client Program uses the SLP client library to make requests. The SLP client library either multicasts requests to SAs or unicasts them to DAs. This communication is transparent to the application except that the unicast method of issuing requests is faster. The behavior of the client library can be affected by setting various SLP configuration properties. For further information, see Chapter 20. The `slpd` process handles all SA functionality, such as answering multicast requests and registering with DAs.




FIGURE 18-3 SLP Implementation

Other SLP Information Sources

Refer to the following documents for further information on SLP:

- Kempf, James, and Pete St. Pierre. *Service Location Protocol for Enterprise Networks*. John Wiley & Sons, Inc. ISBN Number:0-47-3158-7.
- *Authentication Management Infrastructure Administration Guide*. Part Number: 805-1139-03.
- Guttman, Erik, Charles Perkins, John Veizades, and Michael Day. *Service Location Protocol, Version 2, RFC 2608* from the Internet Engineering Task Force (IETF). [<http://www.ietf.org/rfc/rfc2608.txt>]
- Kempf, James, and Erik Guttman. *An API for Service Location, RFC 2614* from the Internet Engineering Task Force (IETF). [<http://www.ietf.org/rfc/rfc2614.txt>]

Composed September 5, 2002

Planning and Enabling SLP (Tasks)

This chapter provides information on planning and enabling SLP. The following sections discuss SLP configuration and the process for enabling SLP.

- “SLP Configuration Considerations” on page 259
- “Using `snoop` to Monitor SLP Activity” on page 260
- “Enabling SLP” on page 263

SLP Configuration Considerations

The SLP daemon is preconfigured with default properties for installation with the Solaris 9 operating environment. If your enterprise functions well with default settings, the SLP deployment requires virtually no administration.

In some situations, however, you might want to modify the SLP properties to tune network operations or to activate certain features. With a few configuration changes you can enable SLP logging, for example. The information in a SLP log and in `snoop` traces can then help you decide if additional configuration is necessary.

SLP configuration properties reside in the `slp.conf` file, which is located in the `/etc/inet` directory. If you decide to change the default property settings, refer to Chapter 20 for the appropriate procedures.

Before you modify SLP configuration settings, consider the following questions that are related to key aspects of network administration:

- What network technologies are operating in the enterprise?
- How much network traffic can the technologies handle smoothly?
- How many services, of what type, are available on the network?
- How many users are on the network? What services do they require? Where are users located in relation to their most frequently accessed services?

Deciding What to Reconfigure

You can use the SLP-enabled `snoop` utility and SLP logging utilities to decide if reconfiguration is necessary and what properties you need to modify. For example, you might reconfigure certain properties to do the following:

- Accommodate a mix of network media that have varying latencies and bandwidth characteristics
 - Recover the enterprise from network failures or unplanned partitioning
 - Add DAs to reduce proliferation of SLP multicasts
 - Implement new scopes to organize users with their most frequently accessed services
-

Using `snoop` to Monitor SLP Activity

The `snoop` utility is a passive administrative tool that provides network traffic information. The utility itself generates minimal traffic and enables you to watch all activity on your network as it occurs.

The `snoop` utility provides traces of the actual SLP message traffic. For example, when you run `snoop` with the `slp` command-line argument, the utility displays traces with information on SLP registrations and deregistrations. You can use the information to gauge the network load by checking which services are being registered and how much reregistration activity is occurring.

The `snoop` utility is also useful for observing the traffic flow between SLP hosts in your enterprise. When you run `snoop` with the `slp` command-line argument, you can monitor the following types of SLP activity to determine if network or agent reconfiguration is needed:

- The number of hosts that are using a particular DA. Use this information to decide whether to deploy additional DAs for load balancing.
- The number of hosts that are using a particular DA. Use this information to help you determine whether to configure certain hosts with new or different scopes.
- Whether UA requests a timeout or DA acknowledgment is slow. You can determine whether a DA is overloaded by monitoring UA timeouts and retransmissions. You can also check if the DA requires more than a few seconds to send registration acknowledgments to an SA. Use this information to rebalance the network load on the DA, if necessary, by deploying additional DAs or changing the scope configurations.

Using `snoop` with the `-v` (verbose) command-line argument, you can obtain registration lifetimes and value of the `fresh` flag in `SrvReg` to determine whether the number of reregistrations should be reduced.

You can also use `snoop` to trace other kinds of SLP traffic, such as the following:

- Traffic between UA clients and DAs
- Traffic between multicasting UA clients and replying SAs

For more information about `snoop`, refer to the `snoop(1M)`.

Tip – Use the `netstat` command in conjunction with `snoop` to view traffic and congestion statistics. For more information about `netstat`, refer to `netstat(1M)`.

▼ How to Use `snoop` to Run SLP Traces

1. **Become superuser.**
2. **Run `snoop` with the `slp` command-line argument.**

Brief Mode:

```
# snoop slp
```

When you run `snoop` in the default *brief* mode, ongoing output is delivered to your screen. SLP messages are truncated to fit on one line per SLP trace.

Verbose Mode:

```
# snoop -v slp
```

When you run `snoop` in *verbose* mode, `snoop` delivers ongoing, unabbreviated output to your screen, which provides the following information:

- The complete address of the service URL
- All service attributes
- The registration lifetime
- All security parameters and flags, if any are available

Note – You can use the `slp` command-line argument with other `snoop` options.

Analyzing a `snoop slp` Trace

In the following example, `slpd` runs on `slphost1` in the default mode as an SA server. The SLP daemon initializes and registers `slphost2` as an echo server. Then, the `snoop slp` process is invoked on `slphost1`.

Note – To simplify the description of the trace results, the lines in the following snoop output are flagged with line numbers.

```
1slphost1 -> 239.255.255.253 SLP V@ SrvRqst [24487] service:directory-agent []
2slphost2 -> slphost1 SLP V2 DAAdvert [24487] service:directory-agent://129
3slphost1 -> 239.255.255.253 SLP V2 SrvRqst [24487] service:directory-agent []
4slphost1 -> 239.255.255.253 SLP V2 SrvRqst [24487] service:directory-agent []
5slphost1 -> slphost2 SLP V2 SrvReg [24488/tcp] service:echo.sun:tcp://slphost1:
6slphost2 -> slphost1 SLP V2 SrvAck [24488/tcp] ok
7slphost1 -> slphost2 SLP V2 SrvDereg [24489/tcp] service:echo.sun:tcp://slphost1:
8slphost2 -> slphost1 SLP V2 SrvAck [24489/tcp] ok
```

1. Shows `slpd` on *slphost1* performing active directory agent discovery by multicasting to the SLP multicast group address in search of directory agents. The message number, 24487, for the active discovery is indicated in square brackets in the trace display.
2. Indicates that the active discovery request 24487 from trace 1 is answered by `slpd`, which is running as a DA on the host *slphost2*. The service URL from *slphost2* has been truncated to fit on a single line. The DA has sent a DA advertisement in reply to the multicast directory agent discovery message, as indicated by the matching message numbers in traces 1 and 2.
3. Shows multicasts from the UAs on *slphost1* for additional DAs. Because *slphost2* has already answered the request, it refrains from responding again, and no other DAs reply.
4. Repeats the multicast operation that is shown in the previous line.
5. Shows a `slpd` on *slphost1* forwarding SA client registrations to the DA on *slphost2*. A unicast service registration (`SrvReg`) for an echo server is made by *slphost1* to the DA on *slphost2*.
6. Shows *slphost2* responding to the *slphost1* `SrvReg` with a service acknowledgment (`SrvAck`) that indicates the registration is successful.

Traffic between the echo server that runs the SA client and the SLP daemon on *slphost1* does not appear in the snoop trace. This absence of information is because the snoop operation is performed over the network loopback.

7. Shows the echo server on *slphost1* deregistering the echo service advertisement. The SLP daemon on *slphost1* forwards the deregistration to the DA on *slphost2*.
8. Shows *slphost2* responding to the *slphost1* with a service acknowledgment (`SrvAck`) that indicates that the deregistration is successful.

The `/tcp` parameter that is appended to the message number on lines 5, 6, 7, and 8 indicates that the message exchange occurred by TCP.

Where to Go From Here

After monitoring the SLP traffic, you can use the information that was collected from the `snoop` traces to help determine whether any reconfiguration of the SLP defaults is needed. Use the related information in Chapter 20 for configuring SLP property settings. For more information about SLP messaging and service registrations, refer to Chapter 22.

Enabling SLP

SLP is enabled by running the SLP daemon, `slpd`. The supported interface for starting `slpd` is the `/etc/init.d/slpd` script, which starts the daemon only if the SLP configuration file, `/etc/inet/slp.conf`, exists. The Solaris operating environment includes the file `/etc/inet/slp.conf.example`. Rename this file to `/etc/inet/slp.conf` to enable SLP at boot time.

Composed September 5, 2002

Administering SLP (Tasks)

The following sections provide information and tasks for configuring SLP agents and processes.

- “Configuring SLP Properties” on page 265
- “Modifying DA Advertising and Discovery Frequency” on page 268
- “Accommodating Different Network Media, Topologies, or Configurations” on page 272
- “Modifying Timeouts on SLP Discovery Requests” on page 277
- “Deploying Scopes” on page 281
- “Deploying DAs” on page 284
- “Multihoming” on page 287

Configuring SLP Properties

SLP configuration properties control network interactions, SLP agent characteristics, status, and logging. In most situations, the default configuration of these properties requires no modification. However, you can use the procedures in this chapter when the network medium or topology changes and to achieve the following goals:

- Compensate for network latencies
- Reduce congestion on the network
- Add agents or reassign IP addresses
- Activate SLP logging

You can edit the SLP configuration file, `/etc/inet/slp.conf`, to perform operations such as those shown in the following table.

TABLE 20-1 SLP Configuration Operations

Operation	Description
Specify whether <code>slpd</code> should act as a DA server. SA server is the default.	Set the <code>net.slp.isDA</code> property to <code>True</code> .
Set timing for DA multicast messages.	Set the <code>net.slp.DAHeartBeat</code> property to control how often a DA multicasts an unsolicited DA advertisement.
Enable DA logging to monitor network traffic.	Set the <code>net.slp.traceDATraffic</code> property to <code>True</code> .

SLP Configuration File: Basic Elements

The `/etc/inet/slp.conf` file defines and activates all SLP activity each time you restart the SLP daemon. The configuration file consists of the following elements:

- Configuration properties
- Comment lines and notations

Configuration Properties

All of the basic SLP properties, such as `net.slp.isDA` and `net.slp.DAHeartBeat`, are named in the following format.

```
net.slp.<keyword>
```

SLP behavior is defined by the value of a property or a combination of properties in the `slp.conf` file. Properties are structured as key-value pairs in the SLP configuration file. As shown in the following example, a key-value pair consists of a property name and an associated setting.

```
<property name>=<value>
```

The key for each property is the property name. The value sets the numeric (distance or time), true/false state, or string value parameters for the property. Property values consist of one of the following data types:

- True/False setting (Boolean)
- Integers
- List of integers
- Strings
- List of strings

Comment Lines and Notations

You can add comments to the `slp.conf` file that describe the nature and function of the line. Comment lines are optional in the file, but can be useful for administration.

Note – Settings in the configuration file are case insensitive. For more information, refer to: Guttman, Erik, James Kempf, and Charles Perkins, “Service Templates and service: scheme,” RFC 2609 from the Internet Engineering Task Force (IETF). [<http://www.ietf.org/rfc/rfc2609.txt>]

▼ How to Change Your SLP Configuration

Use this procedure to change the property settings in your SLP configuration file. SLP-enabled client or service software also can alter the SLP configuration by using the SLP API. This API is documented in “An API for Service Location,” RFC 2614 from the Internet Engineering Task Force (IETF). [<http://www.ietf.org/rfc/rfc2614.txt>]

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Edit the property settings in the `/etc/inet/slp.conf` file as necessary.**
Refer to “Configuration Properties” on page 266 for general information about the SLP property settings. See the sections that follow this procedure for examples of different scenarios in which you might change the `slp.conf` properties. See `slp.conf(4)`.
5. **Save your changes and close the file.**
6. **Restart `slpd` to activate your changes.**

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

Note – The SLP daemon obtains information from the configuration file when you stop or start `slpd`.

For example, you can change the SA server default to enable `slpd` to operate as a DA server by setting the `net.slp.isDA` property to `True` in the `slpd.conf` file.

```
net.slp.isDA=True
```

In each area, various properties control different aspects of the configuration. The following sections describe different scenarios in which you might change the default property settings that are used in SLP configuration.

Modifying DA Advertising and Discovery Frequency

In situations such as the following, you can modify properties that control the timing of DA advertisements and discovery requests .

- When you want the SA or UA to obtain DA configuration information statically from the `net.slp.DAAddresses` property in the `slp.conf` file, you can disable DA discovery.
- When the network is subject to recurrent partitioning, you can change the frequency of passive advertisements and active discovery.
- If UA and SA clients access DAs on the other side of a dial-up connection, you can decrease the DA heartbeat frequency and the active discovery interval to reduce the number of times a dial-up line is activated.
- If network congestion is high, you can limit multicasting.

The procedures in this section explain how to modify the following properties.

TABLE 20-2 DA Advertisement Timing and Discovery Request Properties

Property	Description
<code>net.slp.passiveDADetection</code>	Boolean that specifies whether <code>slpd</code> listens for unsolicited DA advertisements
<code>net.slp.DAActiveDiscoveryInterval</code>	Value that specifies how often <code>slpd</code> performs active DA discovery for a new DA
<code>net.slp.DAHeartBeat</code>	Value that specifies how often a DA multicasts an unsolicited DA advertisement

Limiting UAs and SAs to Statically Configured DAs

Sometimes you might need to limit UAs and SAs to obtaining DA addresses from the static configuration information in the `slp.conf` file. In the next procedure, you can modify two properties that cause `slpd` to obtain DA information exclusively from the `net.slp.DAAddresses` property.

▼ How to Limit UAs and SAs to Statically Configured DAs

Use the following procedure to change the `net.slp.passiveDADetection` and the `net.slp.DAActiveDiscoveryInterval` properties.

Note – Use this procedure only on hosts that execute UAs and SAs which are restricted to static configurations.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Set the `net.slp.passiveDADetection` property to `False` in the `slp.conf` file to disable passive discovery. This setting causes `slpd` to ignore unsolicited DA advertisements.**

```
net.slp.passiveDADetection=False
```
5. **Set the `net.slp.DAActiveDiscoveryInterval` to `-1` to disable initial and periodic active discovery.**

```
net.slp.DAActiveDiscoveryInterval=-1
```
6. **Save your changes and close the file.**
7. **Restart `slpd` to activate your changes.**

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

Configuring DA Discovery for Dial-up Networks

If the UAs or SAs are separated from the DA by a dial-up network, you can configure DA discovery to reduce or eliminate the number of discovery requests and DA advertisements. Dial-up networks usually incur a charge when activated. Minimizing extraneous calls can reduce the cost of using the dial-up network.

Note – You can disable DA discovery completely with the method that is described in “Limiting UAs and SAs to Statically Configured DAs” on page 268.

▼ How to Configure DA Discovery for Dial-up Networks

You can use the following procedure to reduce unsolicited DA advertisements and active discovery by increasing the DA heartbeat period and the active discovery interval.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Increase the `net.slp.DAHeartbeat` property in the `slpd.conf` file.**

```
net.slp.DAHeartbeat=value
```

value

A 32-bit integer that sets the number of seconds for the passive DA advertisement heartbeat

Default Value=10800 seconds (3 hours)

Range of Values=2000–25920000 seconds

For example, you can set the DA heartbeat to approximately 18 hours on a host that is executing a DA:

```
net.slp.DAHeartbeat=65535
```

5. **Increase the `net.slp.DAActiveDiscoveryInterval` property in the `slpd.conf` file:**

```
net.slp.DAActiveDiscoveryInterval value
```

value

A 32-bit integer that sets the number of seconds for DA active discovery queries

Default Value=900 seconds (15 minutes)

Range of Values=300–10800 seconds

For example, you can set the DA active discovery interval to 18 hours on a host that is executing a UA and an SA :

```
net.slp.DAActiveDiscoveryInterval=65535
```

6. **Save your changes and close the file.**

7. **Restart `slpd` to activate your changes.**

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

Configuring the DA Heartbeat for Frequent Partitions

SAs are required to register with all DAs that support their scopes. A DA can appear after `slpd` has performed active discovery. If the DA supports `slpd` scopes, the SLP daemon registers all advertisements on its host with the DA.

One way `slpd` discovers DAs is by the initial unsolicited advertisement a DA sends when it boots. The SLP daemon uses the periodic unsolicited advertisement (the heartbeat) to determine whether a DA is still active. If the heartbeat fails to appear, the daemon removes the DAs the daemon uses and those the daemon offers to UAs.

Finally, when a DA undergoes a controlled shutdown, it transmits a special DA advertisement that informs listening SA services that it will be out of service. The SLP daemon also uses this advertisement to remove inactive DAs from the cache.

If your network is subject to frequent partitions and SAs are long-lived, `slpd` can remove cached DAs during the partitioning if heartbeat advertisements are not received. By decreasing the heartbeat time, you can decrease the delay before a deactivated DA is restored to the cache after the partition is repaired.

▼ How to Configure DA Heartbeat for Frequent Partitions

Use the following procedure to change the `net.slp.DAHeartBeat` property to decrease the DA heartbeat period.

1. **Become superuser.**

2. **Stop `slpd` and all SLP activity on the host.**

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

3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**

4. **Decrease the `net.slp.DAHeartBeat` value to 1 hour (3600 seconds). By default, the DA heartbeat period is set to 3 hours (10800 seconds).**

```
net.slp.DAHeartBeat=3600
```

5. **Save your changes and close the file.**

6. Restart `slpd` to activate your changes.

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

Note – If DA discovery is completely disabled, the `net.slp.DAAddresses` property must be set in `slp.conf` on the hosts that are executing UAs and SAs so that they access the correct DA.

Relieving Network Congestion

If network congestion is high, you can limit the amount of multicast activity. If DAs have not already been deployed in the network, deploying DAs can drastically reduce the amount of SLP-related multicast.

However, even after DAs are deployed, multicast is still necessary for DA discovery. You can reduce the amount of multicast necessary for DA discovery by using the method that is described in “How to Configure DA Discovery for Dial-up Networks” on page 270. You can completely eliminate multicast for DA discovery by using the method that is described in “Limiting UAs and SAs to Statically Configured DAs” on page 268.



Accommodating Different Network Media, Topologies, or Configurations

This section describes possible scenarios in which you can change the following properties to tune SLP performance.

TABLE 20-3 SLP Performance Properties

Property	Description
<code>net.slp.DAAttributes</code>	The minimum refresh interval that a DA accepts for advertisements.
<code>net.slp.multicastTTL</code>	The <i>time-to-live</i> value that is specified for multicast packets.
<code>net.slp.MTU</code>	The byte size set for network packets. The size includes IP and TCP or UDP headers.

TABLE 20-3 SLP Performance Properties (Continued)

Property	Description
<code>net.slp.isBroadcastOnly</code>	The Boolean that is set to indicate if broadcast should be used for DA and non-DA-based service discovery.

Reducing SA Reregistrations

SAs periodically need to refresh their service advertisements before lifetimes expire. If a DA is handling an extremely heavy load from many UAs and SAs, frequent refreshes can cause the DA to become overloaded. If the DA becomes overloaded, UA requests start to time out and are then dropped. UA request timeouts have many possible causes. Before you assume that DA overload is the problem, use a `snoop` trace to check the lifetimes of service advertisements that are registered with a service registration. If the lifetimes are short and reregistrations are occurring often, the timeouts are probably the result of frequent reregistrations.

Note – A service registration is a *reregistration* if the FRESH flag is not set. See Chapter 22 for more information on service registration messages.

▼ How to Reduce SA Reregistrations

Use the following procedure to increase the minimum refresh interval for SAs to reduce reregistrations.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Increase the value of the `min-refresh-interval` attribute of the `net.slp.DAAttributes` property.**

The default minimum reregistration period is zero. The zero default allows SAs to reregister at any point. In the following example, the interval is increased to 3600 seconds (one hour).

```
net.slp.DAAttributes(min-refresh-interval=3600)
```
5. **Save your changes and close the file.**
6. **Restart `slpd` to activate your changes.**

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

Configuring the Multicast Time-to-Live Property

The multicast time-to-live property (`net.slp.multicastTTL`) determines the range over which a multicast packet is propagated on your intranet. The multicast TTL is configured by setting the `net.slp.multicastTTL` property to an integer between 1 and 255. The default value of the multicast TTL is 255, which means, theoretically, that the packet routing is unrestricted. However, a TTL of 255 causes a multicast packet to penetrate the intranet to the border routers on the edge of your administrative domain. Correct configuration of multicast on border routers is required to prevent multicast packets from leaking into the Internet's multicast backbone, or to your ISP.

Multicast TTL scoping is similar to standard IP TTL, with the exception that a TTL comparison is made. Each interface on a router that is multicast enabled is assigned a TTL value. When a multicast packet arrives, the router compares the TTL of the packet with the TTL of the interface. If the TTL of the packet is greater than or equal to the TTL of the interface, the packet TTL is reduced by one, as with the standard IP TTL. If the TTL becomes zero, the packet is discarded. When you use TTL scoping for SLP multicasting, your routers must be properly configured to limit packets to a particular subsection of your intranet.

▼ How to Configure the Multicast Time-to-Live Property

Use the following procedure to reset the `net.slp.multicastTTL` property.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Change the `net.slp.multicastTTL` property in the `slpd.conf` file:**

```
net.slp.multicastTTL=value
```

value

A positive integer less than or equal to 255 that defines the multicast TTL

Note – You can reduce the range of multicast propagation by reducing the TTL value. If the TTL value is 1, then the packet is restricted to the subnet. If the value is 32, the packet is restricted to the site. Unfortunately, the term *site* is not defined by RFC 1075, where multicast TTLs are discussed. Values above 32 refer to theoretical routing on the Internet and should not be used. Values below 32 can be used to restrict multicast to a set of accessible subnets, if the routers are properly configured with TTLs.

5. **Save your changes and close the file.**
6. **Restart `slpd` to activate your changes.**

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

Configuring the Packet Size

The default packet size for SLP is 1400 bytes. The size should be sufficient for most local area networks. For wireless networks or wide area networks, you can reduce the packet size to avoid message fragmentation and reduce network traffic. For local area networks that have larger packets, increasing the packet size can improve performance. You can determine whether the packet size needs to be reduced by checking the minimum packet size for your network. If the network medium has a smaller packet size, you can reduce the `net.slp.MTU` value accordingly.

You can increase the packet size if your network medium has larger packets. However, unless the service advertisements from SAs or queries from UAs frequently overflow the default packet size, you should not have to change the `net.slp.MTU` value. You can use `snoop` to determine whether UA requests often overflow the default packet size and roll over to use TCP rather than UDP.

The `net.slp.MTU` property measures the complete IP packet size, including the link layer header, the IP header, the UDP or TCP header, and the SLP message.

▼ How to Configure the Packet Size

Use the following procedure to change the default packet size by adjusting the `net.slp.MTU` property.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**

4. Change the `net.slp.MTU` property in the `slpd.conf` file:

```
net.slp.MTU=value
```

value

A 16-bit integer that specifies the network packet size, in bytes

Default Value=1400

Range of Values=128–8192

5. Save your changes and close the file.

6. Restart `slpd` to activate your changes.

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

Configuring Broadcast-Only Routing

SLP is designed to use multicast for service discovery in the absence of DAs and for DA discovery. If your network does not deploy multicast routing, you can configure SLP to use broadcast by setting the `net.slp.isBroadcastOnly` property to `True`.

Unlike multicast, broadcast packets do not propagate across subnets by default. For this reason, service discovery without DAs in a non-multicast network works only on a single subnet. In addition, special considerations are required when deploying DAs and scopes on networks in which broadcast is used. A DA on a multihomed host can bridge service discovery between multiple subnets with multicast disabled. See “DA Placement and Scope Name Assignment” on page 291 for more information on deploying DAs on multihomed hosts.

▼ How to Configure Broadcast-Only Routing

Use the following procedure to change `net.slp.isBroadcastOnly` property to `True`.

1. Become superuser.

2. Stop `slpd` and all SLP activity on the host.

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

3. Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.

4. Change the `net.slp.isBroadcastOnly` property in the `slpd.conf` file to `True`:

```
net.slp.isBroadcastOnly=True
```

5. Save your changes and close the file.

6. Restart `slpd` to activate your changes.

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

Modifying Timeouts on SLP Discovery Requests

Two situations might require that you change the timeouts for SLP discovery requests:

- If the SLP agents are separated by multiple subnets, dial-up lines, or other WANs, the network latency can be high enough that the default timeouts are insufficient for a request or registration to be completed. Conversely, if your network is low latency, you can improve performance by decreasing the timeouts.
- If the network is subject to heavy traffic or a high collision rates, the maximum period that SAs and UAs need to wait before sending a message might be insufficient to assure collision-free transactions.

Changing Default Timeouts

High network latency can cause UAs and SAs to time out before a response returns for requests and registrations. Latency can be a problem if a UA is separated from an SA, or if both a UA and an SA are separated from a DA—either by multiple subnets, a dial-up line, or a WAN. You can determine if latency is a problem by checking whether SLP requests are failing because of timeouts on UA and SA requests and registrations. You can also use the `ping` command to measure the actual latency.

The following table lists configuration properties that control timeouts. You can use the procedures in this section to modify these properties.

TABLE 20-4 Time-out Properties

Property	Description
<code>net.slp.multicastTimeouts</code>	The properties that control timeouts for repeated multicast and unicast UDP message transmissions before the transmission is abandoned.
<code>net.slp.DADiscoveryTimeouts</code>	
<code>net.slp.datagramTimeouts</code>	

TABLE 20-4 Time-out Properties (Continued)

Property	Description
net.slp.multicastMaximumWait	The property that controls the maximum amount of time a multicast message is transmitted before it is abandoned.
net.slp.datagramTimeouts	The upper bound of a DA timeout that is specified by the sum of values that are listed for this property. A UDP datagram is repeatedly sent to a DA until a response is received or the time-out bound is reached.

If frequent timeouts are occurring during multicast service discovery or DA discovery, increase the `net.slp.multicastMaximumWait` property from the default value of 15000 milliseconds (15 seconds). Increasing the maximum wait period allows more time for requests on high latency networks to be completed. After you change the `net.slp.multicastMaximumWait`, you should also modify the `net.slp.multicastTimeouts` and `net.slp.DADiscoveryTimeouts`. The sum of the timeout values for these properties equals the `net.slp.multicastMaximumWait` value.

▼ How to Change Default Timeouts

Use the following procedure to change the SLP properties that control timeouts.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Change the `net.slp.multicastMaximumWait` property in the `slpd.conf` file:**

```
net.slp.multicastMaximumWait=value
```

value

A 32-bit integer that lists the sum of the values that are set for
`net.slp.multicastTimeouts` and
`net.slp.DADiscoveryTimeouts`
Default Value=15000 milliseconds (15 seconds)
Range of Values=1000 to 60000 milliseconds

For example, if you determine that multicast requests require 20 seconds (20000 milliseconds), you would adjust the values that are listed for

`net.slp.multicastTimeouts` and the `net.slp.DADiscoveryTimeouts` properties to equal 20000 milliseconds.

```
net.slp.multicastMaximumWait=20000
net.slp.multicastTimeouts=2000,5000,6000,7000
net.slp.DADiscoveryTimeouts=3000,3000,6000,8000
```

5. If necessary, change the `net.slp.datagramTimeouts` property in the `slpd.conf` file:

```
net.slp.datagramTimeouts=value
```

value

A list of 32-bit integers that specify timeouts, in milliseconds, to implement unicast datagram transmission to DAs

Default=3000,3000,3000

For example, you can increase the datagram timeout to 20000 milliseconds to avoid frequent timeouts.

```
net.slp.datagramTimeouts=2000,5000,6000,7000
```

In high-performance networks, you can reduce the time-out bound for multicast and unicast UDP datagram transmission. When you reduce the time-out bound, you decrease latency that is required to satisfy SLP requests.

6. Save your changes and close the file.

7. Restart `slpd` to activate your changes.

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

Configuring the Random-Wait Bound

In networks with heavy traffic or a high collision rate, communication with a DA might be affected. When collision rates are high, the sending agent must retransmit the UDP datagram. You can determine if retransmission is occurring by using `snoop` to monitor traffic on a network of hosts that are running `slpd` as an SA server and a host that is running `slpd` as a DA. If multiple service registration messages for the same service appear in the `snoop` trace from the host that is running `slpd` as an SA server, you might have notice collisions.

Collisions can be particularly troubling at boot time. When a DA first starts, it sends unsolicited advertisements and the SAs respond with registrations. SLP requires the SAs to wait for a random amount of time after receiving a DA advertisement before responding. The random-wait bound is uniformly distributed with a maximum value that is controlled by the `net.slp.randomWaitBound`. The default random-wait bound is 1000 milliseconds (1 second).

▼ How to Configure the Random-Wait Bound

Use the following procedure to change the `net.slp.RandomWaitBound` property in the `slp.conf` file.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Change the `net.slp.RandomWaitBound` property in the `slpd.conf` file:**

```
net.slp.RandomWaitBound=value
```

value

The upper bound for calculating the random-wait time before attempting to contact a DA

Default Value=1000 milliseconds (1 second)

Range of Values=1000 to 3000 milliseconds

For example, you can lengthen the maximum wait to 5000 milliseconds (5 seconds).

```
net.slp.randomWaitBound=5000
```

When you lengthen the random-wait bound, a longer delay in registration occurs. SAs can complete registrations with newly discovered DAs more slowly to avoid collisions and timeouts.

5. **If necessary, change the `net.slp.datagramTimeouts` property in the `slpd.conf` file:**

```
net.slp.datagramTimeouts=value
```

value

A list of 32-bit integers that specify timeouts, in milliseconds, to implement unicast datagram transmission to DAs

Default=3000,3000,3000

For example, you can increase the datagram timeout to 20000 milliseconds to avoid frequent timeouts.

```
net.slp.datagramTimeouts=2000,5000,6000,7000
```

In high-performance networks, you can reduce the time-out bound for multicast and unicast UDP datagram transmission. This setting reduces the amount of latency in satisfying SLP requests.

6. Save your changes and close the file.
7. Restart `slpd` to activate your changes.

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

Deploying Scopes

With scopes, you can provision services that depend on the logical, physical, and administrative groupings of users. You can use scopes to administer access to service advertisements.

Use the `net.slp.useScopes` property to create scopes. For example, in the `/etc/inet/slp.conf` file on a host, add a new scope, called `newscope`, as shown:

```
net.slp.useScopes=newscope
```

Your organization might, for example, have an alcove of networked devices—such as printers and fax machines—at the end of the south hall on the second floor of Building 6. These devices could be used by everyone on the second floor, or you might restrict the usage to members of a certain department. Scopes provide a way for you to provision access to the service advertisements for these machines.

If the devices are dedicated to a single department, you can create a scope with the department name, for example, `mktg`. Devices that belong to other departments can be configured with different scope names.

In another scenario, the departments might be dispersed. For instance, the mechanical engineering and the CAD/CAM departments might be split between floors 1 and 2. However, you can provide the floor 2 machines for the hosts on both floors by assigning them to the same scope. You can deploy scopes in any manner that operates well with your network and users.

Note – UAs that have particular scope are not prevented from actually using services that are advertised in other scopes. Configuring scopes controls only which service advertisements a UA detects. The service is responsible for enforcing any access control restrictions.

When to Configure Scopes

SLP can function adequately without any scope configuration. In the Solaris operating environment, the default scope for SLP is `default`. If no scopes are configured, `default` is the scope of all SLP messages.

You can configure scopes in any of the following circumstances.

- The organizations you support want to restrict service advertisement access to their own members.
- The physical layout of the organization you support suggests that services in a certain area be accessed by particular users.
- The service advertisements that are appropriate for specific users to see must be partitioned.

An example of the first circumstance was cited in “Configuring DA Discovery for Dial-up Networks” on page 269. An example of the second is a situation in which an organization is spread between two buildings, and you want users in a building to access local services in that building. You can configure users in Building 1 with the B1 scope, while you configure users in Building 2 with the B2 scope.

Considerations When Configuring Scopes

When you modify the `net.slp.useScopes` property in the `slpd.conf` file, you configure scopes for all agents on the host. If the host is running any SAs or is acting as a DA, you must configure this property if you want to configure the SAs or DA into scopes other than `default`. If only UAs are running on the machine and the UAs should discover SAs and DAs supporting scopes other than `default`, you do not need to configure the property unless you want to restrict the scopes the UAs use. If the property is not configured, UAs can automatically discover available DAs and scopes through `slpd`. The SLP daemon uses active and passive DA discovery to find DAs, or it uses SA discovery if no DAs are running. Alternatively, if the property is configured, UAs use only the configured scopes and do not discard them.

If you decide to configure scopes, you should consider keeping the `default` scope on the list of configured scopes unless you are sure that all SAs in the network have scopes configured. If any SAs are left unconfigured, UAs with configured scopes are unable to find them. This situation occurs because the unconfigured SAs automatically have scope `default`, but the UAs have the configured scopes.

If you also decide to configure DAs by setting the `net.slp.DAAddresses` property, be sure that the scopes that are supported by the configured DAs are the same as the scopes that you have configured with the `net.slp.useScopes` property. If the scopes are different, `slpd` prints an error message when it is restarted.

▼ How to Configure Scopes

Use the following procedure to add scope names to the `net.slp.useScopes` property in the `slpd.conf` file.

1. Become superuser.

2. Stop `slpd` and all SLP activity on the host.

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

3. Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.

4. Change the `net.slp.useScopes` property in the `slpd.conf` file:

```
net.slp.useScopes=<scope names>
```

scope names

A list of strings that indicate which scopes a DA or SA is allowed to use when making requests, or which scopes a DA must support

Default Value=Default for SA and DA/Unassigned for UA

Note – Use the following to construct scope names:

- Any alphanumeric characters, uppercase or lowercase
- Any punctuation characters (except for: `"`, `\`, `!`, `<`, `=`, `>`, and `~`)
- Spaces that are considered part of the name
- Non-ASCII characters

You use a backslash to escape non-ASCII characters. For example, UTF-8 encoding uses `0xc3a9` hex code to represent the letter *e* with the French *aigue* accent. If the platform does not support UTF-8, you use the UTF-8 hex code as the escape sequence `\c3\a9`.

For example, to specify scopes for `eng` and `mktg` groups in `bldg6`, you change the `net.slp.useScopes` line to the following.

```
net.slp.useScopes=eng,mktg,bldg6
```

5. Save your changes and close the file.

6. Restart `slpd` to activate your changes.

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

Deploying DAs

This section describes the strategic deployment of DAs in a network that is running SLP.

SLP functions adequately with only the base agents (UAs and SAs), and with no deployed DAs or configured scopes. All agents that lack specific configurations use the default scope. DAs serve as caches for service advertisements. Deploying DAs decreases the number of messages that are sent on the network and reduces the time that is required to receive responses to messages. This capability enables SLP to accommodate larger networks.

Why Deploy a SLP DA?

The primary reason to deploy DAs is to reduce the amount of multicast traffic and the delays that are associated with gathering unicast replies. In a large network with many UAs and SAs, the amount of multicast traffic that is involved in service discovery can become so large that network performance degrades. By deploying one or more DAs, UAs must unicast to DAs for service and SAs must register with DAs by using unicast. The only SLP-registered multicast in a network with DAs is for active and passive DA discovery.

SAs register automatically with any DAs they discover within a set of common scopes, rather than accepting multicast service requests. Multicast requests in scopes that are not supported by the DA are still answered directly by the SA, however.

Service requests from UAs are unicast to DAs rather than multicast onto the network when a DA is deployed within the UA's scopes. Consequently, DAs within the UA's scopes reduce multicast. By eliminating multicast for normal UA requests, the time that is required to obtain replies to queries is greatly reduced (from seconds to milliseconds).

DAs act as a focal point for SA and UA activity. Deploying one or several DAs for a collection of scopes provides a centralized point for monitoring SLP activity. By turning on DA logging, it is easier to monitor registrations and requests than by checking the logs from multiple SAs that are scattered around the network. You can deploy any number of DAs for a particular scope or scopes, depending on the need to balance the load.

In networks without multicast routing enabled, you can configure SLP to use broadcast. However, broadcast is very inefficient, because it requires each host to process the message. Broadcast also does not normally propagate across routers. As a result, in a network without multicast routing support, services can be discovered only on the same subnet. Partial support for multicast routing leads to inconsistent ability

to discover services on a network. Multicast messages are used to discover DAs. Partial support for multicast routing, therefore, implies that UAs and SAs register services with all known DAs in the SA's scope. For example, if a UA queries a DA that is called DA1 and the SA has registered services with DA2, the UA will fail to discover a service. See "Configuring Broadcast-Only Routing" on page 276 for more information on how to deploy SLP on networks without multicast enabled.

On a network with inconsistent site-wide support for multicast routing, you must configure the SLP UAs and SAs with a consistent list of DA locations by using the `net.slp.DAAddresses` property.

Finally, the Solaris SLPv2 DA supports interoperability with SLPv1. SLPv1 interoperability is enabled by default in the Solaris DA. If your network contains SLPv1 devices, such as printers, or you need to interoperate with Novell Netware 5, which uses SLPv1 for service discovery, you should deploy a DA. Without a DA, the Solaris SLP UAs are unable to find SLPv1 advertised services.

When to Deploy DAs

Deploy DAs on your enterprise if any of the following conditions are true:

- Multicast SLP traffic exceeds 1 percent of the bandwidth on your network, as measured by `snoop`.
- UA clients experience long delays or timeouts during multicast service requests.
- You want to centralize the monitoring of SLP service advertisements for particular scopes on one or several hosts.
- Your network does not have multicast enabled and consists of multiple subnets that must share services.
- Your network employs devices that support earlier versions of SLP (SLPv1) or you would like SLP service discovery to interoperate with Novell Netware 5.

▼ How to Deploy DAs

Use the following procedure to set the `net.slp.isDA` property to `True` in the `slp.conf` file.

Note – You can assign only one DA per host.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

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

3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**

4. **Set the `net.slp.isDA` property in the `slpd.conf` file to `True`:**

```
net.slp.isDA=True
```

5. **Save your changes and close the file.**

6. **Restart `slpd` to activate your changes.**

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

Where to Place DAs

This section provides suggestions for where to place DAs in different situations.

- When multicast routing is not enabled and DAs are required to bridge service discovery between subnets

In this situation, a DA must be placed on a host with interfaces and all subnets that share services. The `net.slp.interfaces` configuration property does *not* need to be set, unless IP packets are not routed among the interfaces. See “Multihoming Configuration” on page 287 for more information on configuring the `net.slp.interfaces` property.

- When DAs are deployed for scalability and the primary consideration is optimizing agent access

UAs typically make many requests for services to DAs. An SA registers with the DA once, and can refresh the advertisement at periodic but infrequent intervals. As a result, UA access to DAs is far more frequent than SA access. The number of service advertisements is also usually smaller than the number of requests. Consequently, most DA deployments are more efficient if the deployment is optimized for UA access.

- Situating DAs so that they are topologically close to UAs on the network to optimize UA access

Naturally, you must configure the DA with a scope that is shared by both the UA and SA clients.

Placing Multiple DAs for Load Balancing

You can deploy multiple DAs for the same collection of scopes as a means of load balancing. Deploy DAs in any of the following circumstances:

- UA requests to a DA are timing out, or are returning with the `DA_BUSY_NOW` error.
- The DA log shows that many SLP requests are being dropped.
- The network of users who share services in the scopes spans a number of buildings or physical sites.

You can run a `snoop` trace of SLP traffic to determine how many UA requests return with the `DA_BUSY_NOW` error. If the number of UA requests returned is high, UAs in buildings physically and topologically distant from the DA can exhibit slow response or excessive timeouts. In such a scenario, you can deploy a DA in each building to improve response for UA clients within the building.

Links that connect buildings are often slower than the local area networks within the buildings. If your network spans multiple buildings or physical sites, set the `net.slp.DAAddresses` property in the `/etc/inet/slp.conf` file to a list of specific host names or addresses so that the UAs access only the DAs you specify.

If a particular DA is using large amounts of host memory for service registrations, reduce the number of SA registrations by reducing the number of scopes the DA supports. You can split into two a scope that has many registrations. You can then support one of the new scopes by deploying another DA on another host.

Multihoming

A multihomed server acts as a host on multiple IP subnets. The server can sometimes have more than one network interface card and can act as a router. IP packets, including multicast packets, are routed between the interfaces. In some situations, routing between interfaces is disabled. The following sections describe how to configure SLP for such situations.

Multihoming Configuration

Without configuration, `slpd` listens for multicast and for UDP/TCP unicast on the default network interface. If unicast and multicast routing is enabled between interfaces on a multihomed machine, no additional configuration is needed. This is because multicast packets that arrive at another interface are properly routed to the default. As a result, multicast requests for DA or other service advertisements arrive at `slpd`. If routing is not turned on for some reason, configuration is required.

When to Configure for Nonrouted, Multiple Network Interfaces

If one of the following conditions exist, you might need to configure multihomed machines.

- Unicast routing is enabled between the interfaces and multicast routing is disabled.

- Unicast routing and multicast routing are both disabled between the interfaces.

When multicast routing is disabled between interfaces, it is usually because multicast has not been deployed in the network. In that situation, broadcast is normally used for service discovery that is not DA-based and for DA discovery on the individual subnets. Broadcast is configured by setting the `net.slp.isBroadcastOnly` property to `True`.

Configuring Nonrouted, Multiple Network Interfaces (Task Map)

TABLE 20-5 Configuring Nonrouted, Multiple Network Interfaces

Task	Description	For Instructions
Configure the <code>net.slp.interfaces</code> property	Set this property to enable <code>slpd</code> to listen for unicast and multicast/broadcast SLP requests on the specified interfaces.	"Configuring the <code>net.slp.interfaces</code> Property" on page 288
Arrange proxy service advertisements so that UAs on subnets get service URLs with reachable addresses	Restrict proxy advertisement to a machine that is running <code>slpd</code> connected to a single subnet rather than a multihomed host.	"Proxy Advertising on Multihomed Hosts" on page 290
Place the DAs and configure scopes to assure reachability between UAs and SAs	Configure the <code>net.slp.interfaces</code> property on multihomed hosts with a single interface host name or address. Run a DA on a multihomed host, but configure scopes so that SAs and UAs on each subnet use different hosts.	"DA Placement and Scope Name Assignment" on page 291

Configuring the `net.slp.interfaces` Property

If the `net.slp.interfaces` property is set, `slpd` listens for unicast and multicast/broadcast SLP requests on the interfaces that are listed in the property, rather than on the default interface.

Usually, you set the `net.slp.interfaces` property in conjunction with enabling broadcast by setting the `net.slp.isBroadcastOnly` property, because multicast has not been deployed in the network. However, if multicast has been deployed, but is not being routed on this particular multihomed host, a multicast request can arrive at `slpd` from more than one interface. This situation can occur when the routing of packets is handled by another multihomed host or router that connects the subnets that are served by the interfaces.

When such a situation occurs, the SA server or the UA that is sending the request receives two responses from `slpd` on the multihomed host. The responses are then filtered by the client libraries and the client does not see them. The responses are, however, visible in the `snoop` trace.

Note – If unicast routing is turned off, services that are advertised by SA clients on multihomed hosts might not be reachable from all the subnets. If services are unreachable, SA clients can do the following:

- Advertise one service URL for each individual subnet.
 - Assure that requests from a particular subnet are answered with a reachable URL.
-

The SA client library makes no effort to assure that reachable URLs are advertised. The service program, which might or might not handle a multihomed host with no routing, is then responsible for assuring that reachable URLs are advertised.

Before you deploy a service on a multihomed host with unicast routing disabled, use `snoop` to determine whether the service handles requests from multiple subnets correctly. Furthermore, if you plan to deploy a DA on the multihomed host, see “DA Placement and Scope Name Assignment” on page 291.

▼ How to Configure the `net.slp.interfaces` Property

Use the following procedure to change the `net.slp.interfaces` property in the `slp.conf` file.

1. **Become superuser.**
2. **Stop `slpd` and all SLP activity on the host.**

```
# /etc/init.d/slpd stop
```
3. **Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.**
4. **Change the `net.slp.interfaces` property in the `slpd.conf` file:**

```
net.slp.interfaces=value
```

value

List of IPv4 addresses or host names of the network interface cards on which the DA or SA should listen for multicast, unicast UDP, and TCP messages on port 427

For example, a server with three network cards and multicast routing that is turned off is connected to three subnets. The IP addresses of the three network interfaces are

192.147.142.42, 192.147.143.42, and 192.147.144.42. The subnet mask is 255.255.255.0. The following property setting causes `slpd` to listen on all three interfaces for unicast and multicast/broadcast messaging:

```
net.slp.interfaces=192.147.142.42,192.147.143.42,192.147.144.42
```

Note – You can specify IP addresses or resolvable host names for the `net.slp.interfaces` property.

5. Save your changes and close the file.
6. Restart `slpd` to activate your changes.

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

Proxy Advertising on Multihomed Hosts

If a host with multiple interfaces advertises services by using `slpd` and proxy registration, the service URLs that are advertised by `slpd` must contain reachable host names or addresses. If unicast routing is enabled between the interfaces, hosts on all subnets can reach hosts on other subnets. Proxy registrations can also be made for a service on any subnet. If, however, unicast routing is disabled, service clients on one subnet cannot reach services on another subnet through the multihomed host. However, those clients might be able to reach the services through another router.

For example, suppose the host with default host name `bigguy` has three interface cards on three different unrouted subnets. The host names on these subnets are `bigguy`, with IP address 192.147.142.42, `bigguy1`, with IP address 192.147.143.42, and `bigguy2`, with IP address 192.147.144.42. Now suppose that a legacy printer, `oldprinter`, is connected to the 143 subnet and that the URL `service:printing:lpr://oldprinter/queue1` is configured with the `net.slp.interfaces` to listen on all interfaces. The `oldprinter` URL is proxy-advertised on all interfaces. The machines on the 142 and 144 subnets receive the URL in response to service requests, but are unable to access the `oldprinter` service.

The solution to this problem is to perform the proxy advertisement with `slpd` running on a machine that is connected to the 143 subnet only, rather than on the multihomed host. Only hosts on the 143 subnet can obtain the advertisement in response to a service request.

DA Placement and Scope Name Assignment

The placement of DAs and assignment of scope names on a network with a multihomed host must be done carefully to assure that clients obtain accessible services. Be particularly cautious when routing is disabled and the `net.slp.interfaces` property is configured. Again, if unicast routing is enabled between the interfaces on a multihomed machine, no special DA and scope configuration is necessary. The advertisements are cached with the DA identify services that are accessible from any of the subnets. However, if unicast routing is disabled, poor placement of DAs can result in problems.

To see what problems can result in the previous example, consider what would happen if `bigguy` runs a DA, and clients on all subnets have the same scopes. SAs on the 143 subnet register their service advertisements with the DA. UAs on the 144 subnet can obtain those service advertisements, even though hosts on the 143 subnet are unreachable.

One solution to this problem is to run a DA on each subnet and not on the multihomed host. In this situation, the `net.slp.interfaces` property on the multihomed hosts should be configured with a single interface host name or address, or it should be left unconfigured, forcing the default interface to be used. A disadvantage of this solution is that multihomed hosts are often large machines that could better handle the computational load of a DA.

Another solution is to run a DA on the multihomed host, but configure scopes so that the SAs and UAs on each subnet have a different scope. For example, in the previous situation, UAs and SAs on the 142 subnet might have a scope that is called `scope142`. UAs and SAs on the 143 subnet might have another scope that is called `scope143` and UAs and SAs on the 144 subnet could have third scope that is called `scope144`. You can configure the `net.slp.interfaces` property on `bigguy` with the three interfaces so that the DA serves three scopes on the three subnets.

Considerations When Configuring for Nonrouted, Multiple Network Interfaces

Configuring the `net.slp.interfaces` property enables a DA on the multihomed host to bridge service advertisements between the subnets. Such configuration is useful if multicast routing is turned off in the network, but unicast routing between interfaces on a multihomed host is enabled. Because unicast is routed between the interfaces, hosts on a subnet different from the subnet on which the service is located can contact the service when they receive the service URL. Without the DA, SA servers on a particular subnet receive only broadcasts that were made on the same subnet, so they cannot locate services off of their subnet.

Composed September 5, 2002

The most common situation that necessitates configuration of the `net.slp.interfaces` property occurs when multicast is not deployed on the network and broadcast is used instead. Other situations require careful thought and planning to avoid unnecessary duplicate responses or unreachable services.

Incorporating Legacy Services

Legacy services are network services that predate the development and implementation of SLP. Solaris services such as the line printer daemon (lpsched), the NFS file service, and NIS/NIS+ name service, for example, do not contain internal SAs for SLP. This chapter describes when and how to advertise legacy services.

- “When to Advertise Legacy Services” on page 293
- “Advertising Legacy Services” on page 293
- “Considerations When Advertising Legacy Services” on page 297

When to Advertise Legacy Services

With legacy service advertising, you can enable the SLP UAs to find devices and services such as the following on your network:

- Hardware devices and software services that do not contain SLP SAs. When applications with SLP UAs need to find printers or databases that do not contain SLP SAs, for example, legacy advertising might be required.

Advertising Legacy Services

You use any of the following methods to advertise legacy services.

- Modify the service to incorporate an SLP SA.
- Write a small program that advertises on behalf of a service that is not SLP enabled.

- Use proxy advertising to have `slpd` advertise the service.

Modifying the Service

If the source code for the software server is available, you can incorporate a SLP SA. The C and Java APIs for SLP are relatively straightforward to use. See the man pages for information on the C API and documentation on the Java API. If the service is a hardware device, the manufacturer might have an updated PROM that incorporates SLP. Contact the device manufacturer for more information.

Advertising a Service That Is Not SLP Enabled

If the source code or an updated PROM that contains SLP is not available, you can write a small application that uses the SLP client library to advertise the service. This application could function as a small daemon that you start or stop from the same shell script you use to start and stop the service.

SLP Proxy Registration

Solaris `slpd` supports legacy service advertising with a proxy registration file. The proxy registration file is a list of service advertisements in a portable format.

▼ How to Enable SLP Proxy Registration

1. **Create a proxy registration file on the host file system or in any network directory that is accessible by HTTP.**
2. **Determine if a service type template exists for the service.**

The template is a description of the service URL and attributes of a service type. A template is used to define the components of an advertisement for a particular service type:

- If a service type template exists, use the template to construct the proxy registration. See RFC 2609 for more information on service-type templates.
- If a service type template is not available for the service, select a collection of attributes that precisely describe the service. Use a naming authority other than the default for the advertisement. The default naming authority is allowed only for service types that have been standardized. See RFC 2609 for more information on naming authorities.

For example, suppose a company that is called *BizApp* has a local database that is used to track software defects. To advertise the database, the company might use a URL with the service type `service:bugdb.bizapp`. The naming authority would then be `bizapp`.

3. Follow the next steps to configure the `net.slp.serializedRegURL` property in the `/etc/inet/slp.conf` file with the location of the registration file that was created in the previous steps.

4. Become superuser.

5. Stop `slpd` and all SLP activity on the host.

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

6. Back up the default `/etc/inet/slp.conf` file before you change the configuration settings.

7. Specify the location of the proxy registration file in the `net.slp.serializedRegURL` property of the `/etc/inet/slp.conf` file.

```
net.slp.net.slp.serializedRegURL=proxy registration file URL
```

For example, if the serialized registration file is `/net/inet/slp.reg`, you configure the property as shown in the following:

```
net.slp.serializedRegURL=file:/etc/inet/slp.reg
```

8. Save your changes and close the file.

9. Restart `slpd` to activate your changes.

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

Using SLP Proxy Registration to Advertise

A service advertisement consists of lines that identify the service URL, an optional scope, and a series of attribute definitions. The SLP daemon reads, registers, and maintains proxy advertisements exactly as an SA client would. The following is an example of an advertisement from a proxy registration file.

In the example, a legacy printer that supports LPR protocol and an FTP server are advertised. Line numbers have been added for description purposes and are not part of the file.

```
1#Advertise legacy printer.  
2  
3service:lpr://bizserver/mainpool,en,65535  
4scope=eng,corp  
5make-model=Laserwriter II  
6location-description=B16-2345  
7color-supported=monochromatic  
8fonts-supported=Courier,Times,Helvetica 9 10  
9  
10#Advertise FTP server  
11  
12ftp://archive/usr/src/public,en,65535,src-server
```

¹³content=Source code for projects
¹⁴

Note – The proxy registration file supports the same convention for escaping non-ASCII characters as the configuration file does. For more information about the format of the proxy registration file, see RFC 2614.

TABLE 21–1 SLP Proxy Registration File Description

Line Numbers	Description
1 and 10	Comment lines begin with a cross-hatch symbol (#) and do not affect the file’s operation. All characters through the end of a comment line are ignored.
2, 9, and 14	Blank lines that delimit the advertisements.
3, 12	Service URLs that each have three required fields and one optional field that are separated by commas: ■ Generic or service: URL advertised. See RFC 2609 for the specification of how to form a service: URL. ■ Language of the advertisement. In the previous example, the field is designated English, <i>en</i>. Language is an RFC 1766 language tag. ■ Lifetime of the registration, measured in seconds. The lifetime is restricted to an unsigned 16 bit-integer. If the lifetime is less than the maximum, 65535, <i>slpd</i> times out the advertisement. If the lifetime is 65535, <i>slpd</i> refreshes the advertisement periodically, and the lifetime is considered permanent, until <i>slpd</i> exits. ■ (Optional) Service type field – If used, this field defines the service type. If the service URL is defined, you can change the service type under which the URL is advertised. In the previous example of a proxy registration file, line 12 contains a generic FTP URL. The optional type field causes the URL to be advertised under the service type name <i>src-server</i>. The <i>service</i> prefix is not added by default to the type name.
4	Scope designation. Optional line consists of the token <i>scope</i> , followed by an equal sign and a comma-separated list of scope names. Scope names are defined by the <i>net.slp.useScopes</i> configuration property. Only scopes that are configured for the host should be included in the list. When a scope line is not added, the registration is made in all scopes with which <i>slpd</i> is configured. The scope line must appear immediately after the URL line. Otherwise, scope names are recognized as attributes.

TABLE 21-1 SLP Proxy Registration File Description *(Continued)*

Line Numbers	Description
5-8	Attribute definitions. After the optional scope line, the bulk of the service advertisement contains attribute/value list pair lines. Each pair consists of the attribute tag, followed by an equal sign, and an attribute value or a comma-separated list of values. In the previous example of a proxy registration file, line 8 illustrates an attribute list with multiple values. All other lists have single values. The format for the attribute names and values is the same as on-the-wire SLP messages.

Considerations When Advertising Legacy Services

Generally, modifying the source code to add SLP is preferable to writing a SLP-enabled service that uses the SLP API to advertise on behalf of other services. Modifying the source code is also preferable to using proxy registration. When you modify the source code, you can add service-specific features and closely track service availability. If the source code is unavailable, writing an SLP-enabled helper service that advertises on behalf of other services is preferable to using proxy registration. Ideally, this helper service is integrated into the service start/stop procedure that is used to control activation and deactivation. Proxy advertising is generally the third choice, when no source code is available and writing a standalone SA is impractical.

Proxy advertisements are maintained only if `slpd` is running to read the proxy registration file. No direct connection exists between the proxy advertisement and the service. If an advertisement times out or `slpd` is halted, the proxy advertisement is no longer available.

If the service is shut down, `slpd` must be stopped. The serialized registration file is edited to comment out or remove the proxy advertisement, and `slpd` is restarted. You must follow the same procedure when the service is restarted or reinstalled. The lack of connection between the proxy advertisement and the service is a major disadvantage of proxy advertisements.

Composed September 5, 2002

SLP (Reference)

This chapter describes the SLP status codes and message types. SLP message types are listed with the abbreviations and function codes. SLP status codes are shown with descriptions and function codes that are used to indicate that a request is received (code 0), or that the receiver is busy.

Note – The SLP daemon (`slpd`) returns status codes for unicast messages only.

SLP Status Codes

TABLE 22–1 SLP Status Codes

Status Type	Status Code	Description
No Error	0	Request was processed without error.
LANGUAGE_NOT_SUPPORTED	1	For an AttrRqst or SrvRqst, there is data for the service type in the scope, but not in the language that is indicated.
PARSE_ERROR	2	The message fails to follow SLP syntax.
INVALID_REGISTRATION	3	The SrvReg has problems. For example, a zero lifetime or an omitted language tag.

TABLE 22-1 SLP Status Codes *(Continued)*

Status Type	Status Code	Description
SCOPE_NOT_SUPPORTED	4	The SLP message did not include a scope in its scope list that is supported by the SA or DA that answered the request.
AUTHENTICATION_UNKNOWN	5	The DA or SA received a request for an unsupported SLP SPI.
AUTHENTICATION_ABSENT	6	The UA or DA expected URL and attribute authentication in the SrvReg and did not receive it.
AUTHENTICATION_FAILED	7	The UA or DA detected an authentication error in an Authentication block.
VER_NOT_SUPPORTED	9	Unsupported version number in message.
INTERNAL_ERROR	10	An unknown error occurred in the DA or SA. For example, the operating system had no remaining file space.
DA_BUSY_NOW	11	The UA or SA should retry, using exponential backoff. The DA is busy processing other messages.
OPTION_NOT_UNDERSTOOD	12	The DA or SA received an unknown option from the mandatory range.
INVALID_UPDATE	13	The DA received a SrvReg without FRESH set, for an unregistered service or with inconsistent service types.
MSG_NOT_SUPPORTED	14	The SA received an AttrRqst or SrvTypeRqst and does not support it.
REFRESH_REJECTED	15	The SA sent a SrvReg or partial SrvDereg to a DA more frequently than the DA's min-refresh-interval.

SLP Message Types

TABLE 22-2 SLP Message Types

Message Type	Abbreviation	Function Code	Description
Service Request	SrvRqst	1	Issued by a UA to find services or by a UA or SA server during active DA discovery.
Service Reply	SrvRply	2	The DA or SA response to a service request.
Service Registration	SrvReg	3	Enables SAs to register new advertisements, to update existing advertisements with new and changed attributes, and to refresh URL lifetimes.
Service Deregistration	SrvDereg	4	Used by the SA to deregister its advertisements when the service they represent is no longer available.
Acknowledgment	SrvAck	5	The DA response to an SA's service request or service deregistration message.
Attribute Request	AttrRqst	6	Made either by URL or by service type to request a list of attributes.
Attribute Reply	AttrRply	7	Used to return the list of attributes.
DA Advertisement	DAAdvert	8	The DA response to multicast service requests.
Service Type Request	SrvTypeRqst	9	Used to inquire about registered service types that have a particular naming authority and are in a particular set of scopes.
Service Type Reply	SrvTypeRply	10	The message that is returned in response to the service type request.
SA Advertisement	SAAdvert	11	UAs employ the SAAdvert to discover SAs and their scopes in networks where no DAs are deployed.

Composed September 5, 2002

CHAPTER **23**

Mail Services Topics

Chapter 24	Provides overview information for the mail service
Chapter 25	Provides step-by-step instructions for setting up and troubleshooting the mail service
Chapter 26	Provides background information on the mail service
Chapter 27	Provides information on what's new for the mail service

Composed September 5, 2002

Mail Services (Overview)

To set up an electronic mail service and maintain it are complex tasks, both of which are critical to the daily operation of your network. As a network administrator, you need to expand an existing mail service, or perhaps you might need to set up a mail service on a new network or subnet. The chapters on mail services can help you plan and set up a mail service for your network. This chapter provides a list of new features in `sendmail`, as well as a list of other sources of information. The chapter also provides overviews of the software and hardware components that are required to establish a mail service.

- “What’s New in Version 8.12 of `sendmail`” on page 305
- “Introduction to the Components of Mail Services” on page 306

Look in Chapter 25 for procedural information on how to set up and administer mail services. For details, refer to “Mail Services Task Map” on page 310.

Look in Chapter 26 for a more detailed description of the components of mail services. This chapter also describes the mail service programs and files, the mail routing process, and the interactions of `sendmail` with name services.

What’s New in Version 8.12 of `sendmail`

Version 8.12 of `sendmail` has been included in this Solaris 9 release. Chapter 27 describes all of its new features. The following list highlights some of the important changes to `sendmail`.

- “New Configuration File, `submit.cf`” on page 390
- “New or Deprecated Command-Line Options” on page 391
- “New and Revised Configuration File Options and Related Topics” on page 392

- “New and Revised m4 Configuration Macros for `sendmail`” on page 408
- “New Delivery Agent Flags” on page 412
- “New Equates for Delivery Agents” on page 413
- “Changes to Files” on page 417

Chapter 27 also describes these other changes.

- “Changes to `mail.local`” on page 418
- “Changes to `mailstats`” on page 419
- “Changes to `makemap`” on page 419
- “New Command, `editmap`” on page 420
- “Other Changes and Features of Interest” on page 421

Other `sendmail` Information Sources

The following is a list of additional information sources about `sendmail`.

- Costales, Bryan. *sendmail, Second Edition*. O'Reilly & Associates, Inc., 1997.
- Home page for `sendmail` – <http://www.sendmail.org>
- FAQ for `sendmail` – <http://www.sendmail.org/faq>
- README for new `sendmail` configuration files – <http://www.sendmail.org/m4/readme.html>
- A guide for issues that are related to migrating to more recent 8.*+Sun versions of `sendmail` – <http://www.sendmail.org/vendor/sun/>

Introduction to the Components of Mail Services

Many software and hardware components are required to establish a mail service. The following sections give a quick introduction to these components and some of the terms that are used to describe them.

The first section, “Overview of the Software Components” on page 307, defines the terms that are used when discussing the software parts of the mail delivery system. The next section, “Overview of the Hardware Components” on page 307, focuses on the functions of the hardware systems in a mail configuration.

Overview of the Software Components

The following table introduces some of the software components of a mail system. Refer to “Software Components” on page 355 for a complete description of all of the software components.

Component	Description
.forward files	Files that you can set up in a user’s home directory to redirect mail or to send mail to a program automatically
mailbox	A file on a mail server that is the final destination for email messages
mail addresses	Address that contains the name of the recipient and the system to which a mail message is delivered
mail aliases	An alternate name that is used in a mail address
mail queue	A collection of mail messages that needs to be processed by the mail server
postmaster	A special mail alias that is used to report problems and ask questions about the mail service
sendmail configuration file	A file that contains all the information necessary for mail routing

Overview of the Hardware Components

A mail configuration requires three elements, which you can combine on the same system or provide in separate systems.

- A mail host – A system that is configured to handle email addresses that are difficult to resolve
- A minimum of one mail server – A system that is configured to hold one or more mailboxes
- Mail clients – Systems that access mail from a mail server

If users are to communicate with networks outside your domain, you must also add a fourth element, a mail gateway.

Figure 24–1 shows a typical electronic mail configuration, using the three basic mail elements plus a mail gateway.




FIGURE 24-1 Typical Electronic Mail Configuration

Each element is described in detail in "Hardware Components" on page 362.

Mail Services (Tasks)

This chapter describes how to set up and administer mail services. If you are not familiar with administering mail services, read Chapter 24 for an introduction to the components of mail services and for a description of a typical mail service configuration (as shown in Figure 24–1). The following list can help you find groups of related procedures that are covered in this chapter.

- “Mail Services Task Map” on page 310
- “Setting Up Mail Services (Task Map)” on page 314
- “Administering Mail Alias Files (Task Map)” on page 326
- “Administering the Queue Directories (Task Map)” on page 338
- “Administering `.forward` Files (Task Map)” on page 341
- “Troubleshooting Procedures and Tips for Mail Services (Task Map)” on page 344

See Chapter 26 for a more detailed description of the components of mail services. This chapter also describes the mail service programs and files, the mail routing process, and the interactions of `sendmail` with name services.

See Chapter 27 for a description of the new features that are included in version 8.12 of `sendmail`, the version that is in this Solaris 9 release. You can also read about changes to `mail.local`, `mailstats`, and `makemap`. Chapter 27 also provides a description of a new maintenance command, `editmap`.

Mail Services Task Map

The following table refers you to task maps throughout this chapter that focus on a specific group of procedures.

Task	Description	For Instructions
Setting up mail services	Use these procedures to set up each component of your mail service. Learn how to set up a mail server, a mail client, a mail host, a mail gateway, a virtual host, and how to use DNS with <code>sendmail</code> .	"Setting Up Mail Services (Task Map)" on page 314
Building a <code>sendmail</code> configuration file	Use this procedure to modify your <code>sendmail.cf</code> file. See an example of how to enable domain masquerading.	"Building the <code>sendmail.cf</code> Configuration File (Task)" on page 323
Managing mail delivery with an alternate configuration	Use this procedure to prevent mail delivery problems that can occur if the master daemon is disabled.	"Managing Mail Delivery by Using an Alternate Configuration (Task)" on page 325
Administering mail alias files	Use these procedures to provide aliasing on your network. Learn how to manage entries in NIS+ tables. Also, learn how to set up an NIS map, a local mail alias, a keyed map file, and a postmaster alias.	"Administering Mail Alias Files (Task Map)" on page 326
Administering the mail queue	Use these procedures to provide smooth queue processing. Learn how to display and move the mail queue, force mail queue processing, run a subset of the mail queue, and how to run the old mail queue.	"Administering the Queue Directories (Task Map)" on page 338
Administering <code>.forward</code> files	Use these procedures to disable <code>.forward</code> files or change the search path of the <code>.forward</code> file. Also, learn how to permit users to use the <code>.forward</code> file by creating and populating <code>/etc/shells</code> .	"Administering <code>.forward</code> Files (Task Map)" on page 341

Task	Description	For Instructions
Troubleshooting procedures and tips for mail services	Use these procedures and tips to resolve problems with your mail service. Learn how to test the mail configuration, check mail aliases, test the <code>sendmail</code> rule sets, verify connections to other systems, and log messages. Find out where to look for other mail diagnostic information.	"Troubleshooting Procedures and Tips for Mail Services (Task Map)" on page 344
Resolving error messages	Use the information in this section to resolve some mail-related error messages.	"Resolving Error Messages" on page 348

Planning Your Mail System

The following list describes some concerns that should be part of your planning process.

- Determine the type of mail configuration that meets your requirements. This section describes two basic types of mail configuration and briefly lists what you need to set up each configuration. If you need to set up a new mail system or if you are expanding an existing one, you might find this section useful. "Local Mail Only" on page 312 describes the first configuration type, and "Local Mail and a Remote Connection" on page 312 describes the second type.
- As necessary, choose the systems that are to act as mail servers, mail hosts, and mail gateways.
- Make a list of all the mail clients for which you are providing service and include the location of their mailboxes. This list can help you when you are ready to create mail aliases for your users.
- Decide how you plan to update aliases and forward mail messages. You might set up an `aliases` mailbox as a place for users to send requests for mail forwarding and for changes to their default mail alias. If your system uses NIS or NIS+, you can administer mail forwarding, rather than requiring users to manage it themselves. "Administering Mail Alias Files (Task Map)" on page 326 provides a list of tasks that are related to aliasing. "Administering `.forward` Files (Task Map)" on page 341 provides a list of tasks that are related to managing `.forward` files.

After you have completed the planning process, you need to set up systems on your site to perform the functions that are described in "Setting Up Mail Services (Task Map)" on page 314. For other task information, refer to "Mail Services Task Map" on page 310.

Local Mail Only

The simplest mail configuration, shown in Figure 25–1, is two or more workstations that are connected to one mail host. Mail is completely local. All the clients store mail on their local disks and act as mail servers. Mail addresses are parsed by using the `/etc/mail/aliases` files.




FIGURE 25–1 Local Mail Configuration

To set up this kind of mail configuration, you need the following.

- The default `/etc/mail/sendmail.cf` file on each mail client system (no editing is required).
- A server that is designated as the mail host. You can make this designation by adding `mailhost.domain_name` to the `/etc/hosts` file on the mail host. Then, if you are not running NIS or NIS+, add the mail host IP address line to the `/etc/hosts` file of all mail clients.
- Matching `/etc/mail/aliases` files on any system that has a local mailbox (unless you are running NIS or NIS+).
- Enough space in `/var/mail` on each mail client system to hold the mailboxes.

For task information on setting up your mail service, refer to “Setting Up Mail Services (Tasks)” on page 314. If you are looking for a particular procedure that is related to setting up your mail service, refer to “Setting Up Mail Services (Task Map)” on page 314.

Local Mail and a Remote Connection

The most common mail configuration in a small network is shown in Figure 25–2. One system includes the mail server, the mail host, and the mail gateway to the outside world. Mail is distributed by using the `/etc/mail/aliases` files on the mail gateway. No name service is required.




FIGURE 25–2 Local Mail Configuration With a UUCP Connection

In this configuration, you can assume that the mail clients mount their mail files from `/var/mail` on the mail host. To set up this kind of mail configuration, you need the following.

- The `main.cf` file on the mail gateway (no editing is required if Mail Exchanger [MX] records are used).
- The default `/etc/mail/sendmail.cf` file on each mail client system (no editing is required).
- A server that is designated as the mail host. You can make this designation by adding `mailhost.domain_name` to the `/etc/hosts` file on the mail host. Then, if you are not running NIS or NIS+, add the IP address line for the mail host to the `/etc/hosts` file of every mail client.
- Matching `/etc/mail/aliases` files on any system that has a local mailbox (unless you are running NIS or NIS+).
- Enough space in `/var/mail` on the mail server to hold the client mailboxes.

For task information on setting up your mail service, refer to “Setting Up Mail Services (Tasks)” on page 314. If you are looking for a particular procedure that is related to setting up your mail service, refer to “Setting Up Mail Services (Task Map)” on page 314.

Setting Up Mail Services (Task Map)

The following table describes the procedures for setting up mail services.

Task	Description	For Instructions
Setting up a mail server	Steps to enable a server to route mail	"How to Set Up a Mail Server" on page 315
Setting up a mail client	Steps to enable a user to receive mail	"How to Set Up a Mail Client" on page 317
Setting up a mail host	Steps to establish a mail host that can resolve email addresses	"How to Set Up a Mail Host" on page 319
Setting up a mail gateway	Steps to manage communication with networks outside your domain	"How to Set Up a Mail Gateway" on page 320
Using DNS with <code>sendmail</code>	Steps to enable DNS host lookups	"How to Use DNS With <code>sendmail</code> " on page 322
Setting up a virtual host	Steps to assign more than one IP address to a host	"How to Set Up a Virtual Host" on page 322

Setting Up Mail Services (Tasks)

You can readily set up a mail service if your site does not provide connections to email services outside your company or if your company is in a single domain.

Mail requires two types of configurations for local mail. Refer to Figure 25–1 in "Local Mail Only" on page 312 for a representation of these configurations. Mail requires two more configurations for communication with networks outside your domain. Refer to Figure 24–1 in "Overview of the Hardware Components" on page 307 or Figure 25–2 in "Local Mail and a Remote Connection" on page 312 for a representation of these configurations. You can combine these configurations on the same system or provide them on separate systems. For example, if your mail host and mail server functions are on the same system, follow the directions in this section for setting up that system as a mail host. Then, follow the directions in this section for setting up the same system as a mail server.

Note – The following procedures for setting up a mail server and mail client apply when mailboxes are NFS mounted. However, mailboxes typically are maintained in locally mounted `/var/mail` directories, which eliminates the need for the following procedures.

▼ How to Set Up a Mail Server

No special steps are required to set up a mail server that is only serving mail for local users. The user must have an entry in the password file or in the name space, and the user should have a local home directory (for checking the `~/ .forward` file) for mail to be delivered. For this reason, home directory servers are often set up as the mail server. “Hardware Components” on page 362 in Chapter 26 provides more information about the mail server.

The mail server can route mail for many mail clients. The only resource requirement for this type of mail server is that it have adequate spooling space for client mailboxes.

Note – Either the `/var/mail` directory should be available for remote mounting or a service such as Post Office Protocol (POP) or Internet Message Access Protocol (IMAP) should be available from the server for clients to access their mailboxes. The following task shows you how to set up a mail server by using the `/var/mail` directory. To provide configuration guidelines for POP or IMAP is beyond the scope of this document.

For the following task, ensure that the `/etc/dfs/dfstab` file shows that the `/var/mail` directory is exported.

1. **Become superuser on the mail server or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Stop `sendmail`.**

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

3. **Check if the `/var/mail` directory is available for remote access.**

```
# share
```

If the `/var/mail` directory is listed, proceed to step 5.

If the `/var/mail` directory is not listed or if no list appears, continue with the appropriate substep.

- a. **(Optional) If no list appears, start NFS services.**

Follow the procedure, “How to Set Up Automatic File-System Sharing” on page 152, to use the /var/mail directory to start NFS services.

- b. (Optional) If the /var/mail directory is not included in the list, add it to /etc/dfs/dfstab.

Add the following command line to the /etc/dfs/dfstab file.

```
share -F nfs -o rw /var/mail
```

- 4. Make the file system available for mounting.

```
# shareall
```

- 5. Ensure that your name service has been started.

- a. (Optional) If you are running NIS, use this command.

```
# ypwhich
```

For more information, refer to the ypwhich(1) man page.

- b. (Optional) If you are running NIS+, use this command.

```
# nisl
```

For more information, refer to the nisl(1) man page.

- c. (Optional) If you are running DNS, use this command.

```
# nslookup hostname
```

hostname Use your host name.

For more information, refer to the nslookup(1M) man page.

- d. (Optional) If you are running LDAP, use this command.

```
# ldaplist
```

For more information, refer to the ldaplist(1) man page.

- 6. Restart sendmail.

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

Note – The mail.local program automatically creates mailboxes in the /var/mail directory the first time a message is delivered. You do not need to create individual mailboxes for your mail clients.

▼ How to Set Up a Mail Client

A mail client is a user of mail services with a mailbox on a mail server and a mail alias in the `/etc/mail/aliases` file that points to the location of the mailbox. “Hardware Components” on page 362 in Chapter 26 provides a brief description of a mail client.

Note – You can also perform the task of setting up a mail client by using a service such as Post Office Protocol (POP) or Internet Message Access Protocol (IMAP). However, to provide configuration guidelines for POP or IMAP is beyond the scope of this document.

- 1. Become superuser on the mail client’s system or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

- 2. Stop `sendmail`.**

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

- 3. Ensure that a `/var/mail` mount point exists on the mail client’s system.**

The mount point should have been created during the installation process. You can use `ls` to ensure that the file system exists. The following example shows the response you receive if the file system has not been created.

```
# ls -l /var/mail
/var/mail not found
```

- 4. Ensure that no files are in the `/var/mail` directory.**

If mail files do exist in this directory, you should move them so that they are not covered when the `/var/mail` directory is mounted from the server.

- 5. Mount the `/var/mail` directory from the mail server.**

You can mount the mail directory automatically or at boot time.

- a. (Optional) Mount `/var/mail` automatically.**

Add an entry such as the following to the `/etc/auto_direct` file.

```
/var/mail -rw,hard,actimeo=0 server:/var/mail
```

server

Use the assigned server name.

- b. (Optional) Mount `/var/mail` at boot time.**

Add the following entry to the `/etc/vfstab` file. This entry permits the `/var/mail` directory on the mail server that is specified to mount the local `/var/mail` directory.

```
server:/var/mail - /var/mail nfs - no rw,hard,actimeo=0
```

The client's mailbox is automatically mounted any time the system is rebooted. If you are not rebooting the system, type the following command to mount the client mailbox.

```
# mountall
```



Caution – For mailbox locking and mailbox access to work properly, you must include the `actimeo=0` option when mounting mail from an NFS server.

6. Update `/etc/hosts`.

Edit the `/etc/hosts` file and add an entry for the mail server. This step is not required if you are using a name service.

```
# cat /etc/hosts
#
# Internet host table
#
..
IP_address      mailhost mailhost mailhost.example.com
```

IP_address Use the assigned IP addresses.

example.com Use the assigned domain.

mailhost Use the assigned mailhost.

For more information, refer to the `hosts(4)` man page.

7. Add an entry for the client to one of the alias files.

Refer to “Administering Mail Alias Files (Task Map)” on page 326 for a task map about administering mail alias files.

Note – The `mail.local` program automatically creates mailboxes in the `/var/mail` directory the first time a message is delivered. You do not need to create individual mailboxes for your mail clients.

8. Restart `sendmail`.

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

▼ How to Set Up a Mail Host

A mail host resolves email addresses and reroutes mail within your domain. A good candidate for a mail host is a system that connects your network to the outside world or to a parent domain. The following procedure shows you how to set up a mail host.

1. **Become superuser on the mail host system or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Stop sendmail.**

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

3. **Verify the host-name configuration.**

Run the `check-hostname` script to verify that `sendmail` can identify the fully qualified host name for this server.

```
% /usr/lib/mail/sh/check-hostname
hostname phoenix OK: fully qualified as phoenix.example.com
```

If this script is not successful in identifying the fully qualified host name, you need to add the fully qualified host name as the first alias for the host in `/etc/hosts`.

4. **Update the `/etc/hosts` file.**

Choose the step that is appropriate for you.

- a. **(Optional) If you are using NIS or NIS+, edit the `/etc/hosts` file on the system that is assigned to be the new mail host.**

Add the word `mailhost` and `mailhost.domain` after the IP address and system name of the mail host system.

```
IP_address mailhost mailhost mailhost.domain loghost
```

<code>IP_address</code>	Use the assigned IP address.
<code>mailhost</code>	Use the system name of the mail host system.
<code>domain</code>	Use the expanded domain name.

The system is now designated as a mail host. The *domain* should be identical to the string that is given as the subdomain name in the output of the following command.

```
% /usr/lib/sendmail -bt -d0 </dev/null
Version 8.12.0+Sun
Compiled with: LDAPMAP MAP_REGEX LOG MATCHGECOS MIME7TO8 MIME8TO7
               NAMED_BIND NDBM NETINET NETINET6 NETUNIX NEWDB NIS
               NISPLUS QUEUE SCANF SMTP USERDB XDEBUG

===== SYSTEM IDENTITY (after readcf) =====
```

```
(short domain name) $w = phoenix
(canonical domain name) $j = phoenix.example.com
(subdomain name) $m = example.com
(node name) $k = phoenix
```

=====

See the following example of how the `hosts` file should look after these changes.

```
# cat /etc/hosts
#
# Internet host table
#
172.31.255.255    localhost
192.168.255.255  phoenix mailhost mailhost.example.com loghost
```

- b. (Optional) If you are not using NIS or NIS+, edit the `/etc/hosts` file on each system in the network and create the following entry.**

```
IP_address mailhost mailhost mailhost.domain loghost
```

5. Select the correct configuration file to copy and rename.

The following command copies and renames the `/etc/mail/main.cf` file.

```
# cp /etc/mail/main.cf /etc/mail/sendmail.cf
```

6. Restart `sendmail`.

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

7. Test your mail configuration.

See “How to Test the Mail Configuration” on page 344 for instructions.

For further information about mail hosts, refer to “Hardware Components” on page 362 in Chapter 26.

▼ How to Set Up a Mail Gateway

A mail gateway manages communication with networks outside your domain. The mailer on the sending mail gateway can match the mailer on the receiving system.

A good candidate for a mail gateway is a system that is attached to Ethernet and phone lines or a system that is configured as a router to the Internet. You can configure the mail host or another system as the mail gateway. You might choose to configure more than one mail gateway for your domain. If you have UNIX-to-UNIX Copy Program (UUCP) connections, you should configure the system (or systems) with UUCP connections as the mail gateway.

1. Become superuser on the mail gateway or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Stop sendmail.

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

3. Select the correct configuration file to copy and rename.

The following command copies and renames the `main.cf` file.

```
# cp /etc/mail/main.cf /etc/mail/sendmail.cf
```

4. Verify the host-name configuration.

Run the `check-hostname` script to verify that `sendmail` can identify the fully qualified host name for this server.

```
# /usr/lib/mail/sh/check-hostname
hostname phoenix OK: fully qualified as phoenix.example.com
```

If this script is not successful in identifying the fully qualified host name, you need to add the fully qualified host name as the first alias for the host in `/etc/hosts`. If you need help with this step, refer to Step 4 of “How to Set Up a Mail Host” on page 319.

5. Ensure that your name service has been started.

a. (Optional) If you are running NIS, use this command.

```
# ypwhich
```

For more information, refer to the `ypwhich(1)` man page.

b. (Optional) If you are running NIS+, use this command.

```
# nislsl
```

For more information, refer to the `nislsl(1)` man page.

c. (Optional) If you are running DNS, use this command.

```
# nslookup hostname
```

hostname

Use your host name.

For more information, refer to the `nslookup(1M)` man page.

d. (Optional) If you are running LDAP, use this command.

```
# ldaplist
```

For more information, refer to the `ldaplist(1)` man page.

6. Restart sendmail.

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

7. Test your mail configuration

See “How to Test the Mail Configuration” on page 344 for instructions.

For more information about the mail gateway, refer to “Hardware Components” on page 362 in Chapter 26.

▼ How to Use DNS With `sendmail`

The DNS name service does not support aliases for individuals. This name service does support aliases for hosts or domains that use Mail Exchange (MX) records and cname records. You can specify host names, domain names, or both names in the DNS database. For more information about `sendmail` and DNS, see “Interactions of `sendmail` With Name Services” on page 383 in Chapter 26, or see the *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

1. Become superuser or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Enable DNS host lookups (NIS+ only).

Edit the `/etc/nsswitch.conf` file and remove the `#` from the `hosts` definition that includes the `dns` flag. The host entry must include the `dns` flag, as the following example shows, in order for the DNS host aliases to be used.

```
# grep hosts /etc/nsswitch.conf
#hosts:      nisplus [NOTFOUND=return] files
hosts:      nisplus dns [NOTFOUND=return] files
```

3. Check for a `mailhost` and `mailhost.domain` entry.

Use `nslookup` to ensure that an entry exists for `mailhost` and `mailhost.domain` in the DNS database. For more information, refer to the `nslookup(1M)` man page.

▼ How to Set Up a Virtual Host

If you need to assign more than one IP address to a host, see this Web site: <http://www.sendmail.org/virtual-hosting.html>. This site provides complete instructions on how to use `sendmail` to set up a virtual host. However, in the “Sendmail Configuration” section, do not perform step 3b, as shown in the following.

```
# cd sendmail-VERSION/cf/cf
# ./Build mailserver.cf
# cp mailserver.cf /etc/mail/sendmail.cf
```

Instead, for the Solaris operating environment, perform the following steps.

```
# cd /usr/lib/mail/cf
# /usr/ccs/bin/make mailserver.cf
# cp mailserver.cf /etc/mail/sendmail.cf
```

mailserver

Use the name of the `.cf` file.

“Building the `sendmail.cf` Configuration File (Task)” on page 323 outlines these same three steps as part of the build process.

After you have generated your `/etc/mail/sendmail.cf` file, you can continue with the next steps to create a virtual user table, and so forth.

Building the `sendmail.cf` Configuration File (Task)

“How to Build a New `sendmail.cf` File” on page 323 shows you how to build the configuration file. Although you can still use older versions of `sendmail.cf` files, the best practice is to use the new format.

For more details, you should read from the following resources.

- `/usr/lib/mail/README` provides a complete description of the configuration process.
- <http://www.sendmail.org> provides online information about `sendmail` configuration.
- “Versions of the Configuration File” on page 354 and “`sendmail` Configuration File” on page 375, in Chapter 26, also provide some guidance.

The following sections in Chapter 27 identify new `m4` configuration features.

- “New and Revised Configuration File Options and Related Topics” on page 392
- “New and Revised `m4` Configuration Macros for `sendmail`” on page 408

▼ How to Build a New `sendmail.cf` File

The following procedure shows you how to build a new configuration file.

Note – `/usr/lib/mail/cf/main-v7sun.mc` is now `/usr/lib/mail/cf/main.mc`.

1. Become superuser or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Stop sendmail.

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

3. Make a copy of the configuration files that you are changing.

```
# cd /usr/lib/mail/cf
# cp main.mc myhost.mc
```

myhost Select a new name for your .mc file.

4. Edit the new configuration files (for example, *myhost.mc*), as necessary.

For example, add the following command line to enable domain masquerading.

```
# cat myhost.mc
..
MASQUERADE_AS( 'host.domain' )
```

host.domain Use the desired host name and domain name.

In this example, MASQUERADE_AS causes mail that is sent to be labeled as originating from *host.domain*, rather than \$j.

5. Build the configuration file by using `m4`.

```
# /usr/ccs/bin/make myhost.cf
```

6. Test the new configuration file by using the `-C` option to specify the new file.

```
# /usr/lib/sendmail -C myhost.cf -v testaddr </dev/null
```

This command sends a message to `testaddr` while displaying messages as it runs. Only outgoing mail can be tested without restarting the `sendmail` service on the system. For systems that are not handling mail yet, use the full testing procedure in “How to Test the Mail Configuration” on page 344.

7. Install the new configuration file after making a copy of the original.

```
# cp /etc/mail/sendmail.cf /etc/mail/sendmail.cf.save
# cp myhost.cf /etc/mail/sendmail.cf
```

8. Restart the `sendmail` service.

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

Managing Mail Delivery by Using an Alternate Configuration (Task)

To facilitate the transport of inbound and outbound mail, the new default configuration of `sendmail` uses a daemon and a client queue runner. If you have disabled your daemon, you should perform the following task. For a detailed explanation, refer to “New Configuration File, `submit.cf`” on page 390.

▼ How to Manage Mail Delivery by Using an Alternate Configuration of `sendmail.cf`

In the default configuration of `sendmail`, the client queue runner must be able to submit mail to the daemon on the local SMTP port. If the daemon is not listening on the SMTP port, the mail remains in the queue. To avoid this problem, perform the following task. For more information about the daemon and client queue runner and to understand why you might have to use this alternate configuration, refer to “New Configuration File, `submit.cf`” on page 390.

This procedure ensures that your daemon runs only to accept connections from the local host.

1. Become superuser or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Stop `sendmail`.

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

3. Make a copy of the configuration file (either `subsidiary.mc` or `main.mc`, depending on your requirements) that you are changing. In this example, the `subsidiary.mc` file is used.

```
# cd /usr/lib/mail/cf
# cp subsidiary.mc myhost.mc
```

myhost

Select a new name for your `.mc` file.

4. Edit the new configuration file (for example, `myhost.mc`).

Add the following lines before the `MAILER()` lines.

```
# cat myhost.mc
..
```

```
DAEMON_OPTIONS('NAME=NoMTA4, Family=inet, Addr=127.0.0.1')dnl
DAEMON_OPTIONS('NAME=NoMTA6, Family=inet6, Addr>:::1')dnl
```

5. Build the configuration file by using `m4`.

```
# /usr/ccs/bin/make myhost.cf
```

6. Install the new configuration file after making a copy of the original.

```
# cp /etc/mail/sendmail.cf /etc/mail/sendmail.cf.save
# cp myhost.cf /etc/mail/sendmail.cf
```

7. Restart the `sendmail` service.

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

Administering Mail Alias Files (Task Map)

The following table describes the procedures for administering mail alias files. For more information on this topic, refer to “Mail Alias Files” on page 376 in Chapter 26.

Task	Description	For Instructions
Managing alias entries in an NIS+ <code>mail_aliases</code> table	If your name service is NIS+, use these procedures to manage the contents of your <code>mail_aliases</code> table. Learn how to list, add, edit, and delete entries.	“How to Manage Alias Entries in an NIS+ <code>mail_aliases</code> Table” on page 327
Setting up an NIS <code>mail_aliases</code> map	If your name service is NIS, follow these instructions to facilitate aliasing with a <code>mail_aliases</code> map.	“How to Set Up an NIS <code>mail_aliases</code> Map” on page 332
Setting up a local mail alias file	If you are not using a name service (such as NIS or NIS+), follow these instructions to facilitate aliasing with the <code>/etc/mail/aliases</code> file.	“How to Set Up a Local Mail Alias File” on page 333
Creating a keyed map file	Use these steps to facilitate aliasing with a keyed map file.	“How to Create a Keyed Map File” on page 335
Setting up the postmaster alias	Use the procedures in this section to manage the postmaster alias. You must have this alias.	“Managing the postmaster Alias” on page 335

Administering Mail Alias Files (Tasks)

Mail aliases must be unique within the domain. This section provides the procedures for administering mail alias files. Alternately, you can use the AdminTool's Database Manager application to perform these tasks on the aliases database.

In addition, you can create database files for the local mail host by using `makemap`. Refer to the `makemap(1M)` man page. The use of these database files does not provide all of the advantages of using a name service such as NIS or NIS+. However, you should be able to retrieve the data from these local database files faster because no network lookups are involved. For more information, refer to "Interactions of `sendmail` With Name Services" on page 383 and "Mail Alias Files" on page 376 in Chapter 26.

▼ How to Manage Alias Entries in an NIS+ `mail_aliases` Table

To manage entries in an NIS+ table, you can use the `aliasadm` command. To list, add, modify, or delete table entries with the `aliasadm` command, you begin a particular task with the following steps.

1. **Either be a member of the NIS+ group that owns the table, or become `root` on the mail server, or assume an equivalent role.**
For information about roles, refer to "Using Privileged Applications" in *System Administration Guide: Security Services*.
2. **Complete your task by following the instructions from the example that meets your requirements.**
 - "Example—Initiating an NIS+ `mail_aliases` Table" on page 328
 - "Example—Listing the Entire Contents of the NIS+ `mail_aliases` Table" on page 328
 - "Example—Listing an Individual Entry From the NIS+ `mail_aliases` Table" on page 328
 - "Example—Listing Partial Matches From the NIS+ `mail_aliases` Table" on page 329
 - "Example—Adding Aliases to the NIS+ `mail_aliases` Table From the Command Line" on page 329
 - "Example—Adding Entries by Editing an NIS+ `mail_aliases` Table" on page 330
 - "Example—Editing Entries in an NIS+ `mail_aliases` Table" on page 331
 - "Example—Deleting Entries From an NIS+ `mail_aliases` Table" on page 332

In some instances, you should begin the task by compiling a list of each of your mail clients, the locations of their mailboxes, and the names of the mail server systems.

▼ Example—Initiating an NIS+ mail_aliases Table

To create a table, follow these instructions.

1. **Either be a member of the NIS+ group that owns the table, or become root on the mail server, or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Initiate an NIS+ table.**

```
# aliasadm -I
```

3. **Add entries to the table.**

- To add two or three aliases, refer to “Example—Adding Aliases to the NIS+ mail_aliases Table From the Command Line” on page 329.
- To add more than two or three aliases, refer to “Example—Adding Entries by Editing an NIS+ mail_aliases Table” on page 330.

For more information, refer to the `aliasadm(1M)` man page.

▼ Example—Listing the Entire Contents of the NIS+ mail_aliases Table

To see a complete list of the contents of the table, follow these instructions.

1. **Either be a member of the NIS+ group that owns the table, or become root on the mail server, or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **List all of the entries in alphabetical order by alias.**

```
# aliasadm -l
```

For more information, refer to the `aliasadm(1M)` man page.

▼ Example—Listing an Individual Entry From the NIS+ mail_aliases Table

To see an individual entry from the table, follow these instructions.

1. **Either be a member of the NIS+ group that owns the table, or become root on the mail server, or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. List an individual entry.

```
# aliasadm -m ignatz
ignatz: ignatz@saturn # Alias for Iggy Ignatz
```

The command matches only the complete alias name, not partial strings. You cannot use metacharacters (such as * and ?) with `aliasadm -m`.

For more information, refer to the `aliasadm(1M)` man page.

▼ Example—Listing Partial Matches From the NIS+ `mail_aliases` Table

To see partial matches from the table, follow these instructions.

1. Either be a member of the NIS+ group that owns the table, or become `root` on the mail server, or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. List partial matches from the table.

```
# aliasadm -l | grep partial_string
```

`partial_string` Use the string of your choice for your search.

For more information, refer to the `aliasadm(1M)` man page.

▼ Example—Adding Aliases to the NIS+ `mail_aliases` Table From the Command Line

To add two or three aliases to the table, follow these instructions.

1. Compile a list of each of your mail clients, the locations of their mailboxes, and the names of the mail server systems.

2. Either be a member of the NIS+ group that owns the table, or become `root` on the mail server, or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

3. (Optional) If necessary, initiate an NIS+ table.

If you are creating a completely new NIS+ `mail_aliases` table, you must first initiate the table. To complete this task, refer to “Example—Initiating an NIS+

mail_aliases Table" on page 328.

4. Add aliases to the table.

See this example of a typical entry.

```
# aliasadm -a iggy iggy.ignatz@saturn "Iggy Ignatz"
```

The following list describes the input from the preceding example.

-a	The option for adding an alias
iggy	The short form of the alias name
iggy.ignatz@saturn	The expanded alias name
"Iggy Ignatz"	The name for the alias in quotation marks

5. Display the entry you created and ensure that it is correct.

```
# aliasadm -m alias
```

alias	The entry that you created

For more information, refer to the aliasadm(1M) man page.

▼ Example—Adding Entries by Editing an NIS+ mail_aliases Table

To add more than two or three aliases to the table, follow these instructions.

1. **Compile a list of each of your mail clients, the locations of their mailboxes, and the names of the mail server systems.**
2. **Either be a member of the NIS+ group that owns the table, or become root on the mail server, or assume an equivalent role.**

For information about roles, refer to "Using Privileged Applications" in *System Administration Guide: Security Services*.

3. Display and edit the aliases table.

```
# aliasadm -e
```

This command displays the table and enables you to edit it. The editor you use has been set with the \$EDITOR environment variable. If this variable is not set, vi is the default editor.

4. Use the following format to type each alias on a separate line.

```
alias: expanded_alias # ["option" # "comments"]
```

<i>alias</i>	This column is for the short form of the alias name.
<i>expanded_alias</i>	This column is for the expanded alias name.
<i>option</i>	This column is reserved for future use.
<i>comments</i>	This column is used for comments about the individual alias, such as a name for the alias.

If you leave the option column blank, type an empty pair of quotation marks (" ") and add the comments.

The order of the entries is not important to the NIS+ `mail_aliases` table. The `aliasadm -l` command sorts the list and displays the entries in alphabetical order. For more information, refer to “Mail Alias Files” on page 376 and the `aliasadm(1M)` man page.

▼ Example—Editing Entries in an NIS+ `mail_aliases` Table

To edit entries in the table, follow these instructions.

1. **Either be a member of the NIS+ group that owns the table, or become `root` on the mail server, or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Display the alias entry.**

```
# aliasadm -m alias
```

<i>alias</i>	Use the assigned alias name.

3. **Edit the alias entry, as necessary.**

```
# aliasadm -c alias expanded_alias [options comments]
```

<i>alias</i>	If necessary, edit the alias name.
<i>expanded_alias</i>	If necessary, edit the expanded alias name.
<i>options</i>	If necessary, edit the option.
<i>comments</i>	If necessary, edit the comment for this entry.

For more information, refer to the `aliasadm(1M)` man page, as well as “Mail Alias Files” on page 376.

4. Display the entry that you have edited and ensure that the entry is correct.

```
# aliasadm -m alias
```

For more information, refer to the `aliasadm(1M)` man page.

▼ Example—Deleting Entries From an NIS+ `mail_aliases` Table

To delete entries from the table, follow these instructions.

1. **Either be a member of the NIS+ group that owns the table, or become `root` on the mail server, or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Delete an entry from the table.**

```
# aliasadm -d alias
```

alias

Use the alias name for the entry that you are deleting.

For more information, refer to the `aliasadm(1M)` man page.

▼ How to Set Up an NIS `mail_aliases` Map

Use the following procedure to facilitate aliasing with an NIS `mail_aliases` map.

1. **Compile a list of each of your mail clients, the locations of their mailboxes, and the names of the mail server systems.**

2. **Become `root` on the NIS master server or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

3. **Edit the `/etc/mail/aliases` file, and make the following entries.**

- a. **Add an entry for each mail client.**

```
# cat /etc/mail/aliases
..
alias:expanded_alias
```

Composed September 5, 2002

<i>alias</i>	Use the short alias name.
<i>expanded_alias</i>	Use the expanded alias name (<i>user@host.domain.com</i>).

b. Ensure that you have a `Postmaster: root` entry.

```
# cat /etc/mail/aliases
..
Postmaster: root
```

c. Add an alias for `root`. Use the mail address of the person who is designated as the postmaster.

```
# cat /etc/mail/aliases
..
root: user@host.domain.com
```

<i>user@host.domain.com</i>	Use the assigned address of the designated postmaster.

4. Ensure that the NIS master server is running a name service to resolve the host names on each mail server.

5. Change to the `/var/yp` directory.

```
# cd /var/yp
```

6. Apply the `make` command.

```
# make
```

The changes in the `/etc/hosts` and `/etc/mail/aliases` files are propagated to NIS slave systems and are active in only a few minutes, at most.

▼ How to Set Up a Local Mail Alias File

Use the following procedure to resolve aliases with a local mail alias file.

1. Compile a list of each of your users and the locations of their mailboxes.

2. Become `root` on the mail server or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

3. Edit the `/etc/mail/aliases` file and make the following entries.

a. Add an entry for each user.

```
user1: user2@host.domain
```

<i>user1</i>	Use the new alias name.
<i>user2@host.domain</i>	Use the actual address for the new alias.

b. Ensure that you have a `Postmaster: root` entry.

```
# cat /etc/mail/aliases
..  
Postmaster: root
```

c. Add an alias for `root`. Use the mail address of the person who is designated as the postmaster.

```
# cat /etc/mail/aliases
..  
root: user@host.domain.com
```

<i>user@host.domain.com</i>	Use the assigned address of the designated postmaster.

4. Rebuild the alias database.

```
# newaliases
```

Depending on the configuration of the `AliasFile` option in `/etc/mail/sendmail.cf`, this command generates in binary form either the single file, `/etc/mail/aliases.db`, or the pair of files, `/etc/mail/aliases.dir` and `/etc/mail/aliases.pag`.

5. Perform one of the following steps to copy the file or files that were generated.

a. (Optional) Copy the `/etc/mail/aliases`, the `/etc/mail/aliases.dir`, and the `/etc/mail/aliases.pag` files to each of the other systems.

You can copy the three files by using the `rcp` or `rdist` commands. Refer to the `rcp(1)` man page or the `rdist(1)` man page for more information. Alternately, you can create a script for this purpose.

When you copy these files, you do not need to run the `newaliases` command on each of the other systems. However, remember that you must update all the `/etc/mail/aliases` files each time you add or remove a mail client.

b. (Optional) Copy the `/etc/mail/aliases.db` file to each of the other systems.

You can copy the file by using the `rcp` or `rdist` commands. Refer to the `rcp(1)` man page or the `rdist(1)` man page for more information. Alternately, you can create a script for this purpose.

When you copy this file, you do not need to run the `newaliases` command on each of the other systems. However, remember that you must update all the `/etc/mail/aliases` files each time you add or remove a mail client.

▼ How to Create a Keyed Map File

To create a keyed map file, follow these instructions.

1. Become superuser on the mail server or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Create an input file.

Entries can have the following syntax.

```
old_name@newdomain.com    new_name@newdomain.com
old_name@olddomain.com    error:nouser No such user here
@olddomain.com            %1@newdomain.com
```

<i>old_name@newdomain.com</i>	Use the user name that was previously assigned with the domain that is newly assigned.

<i>new_name@newdomain.com</i>	Use the address that is newly assigned.

<i>old_name@olddomain.com</i>	Use the user name that was previously assigned with the domain that was previously assigned.

<i>olddomain.com</i>	Use the domain that was previously assigned.

<i>newdomain.com</i>	Use the domain that is newly assigned.

The first entry redirects mail to a new alias. The next entry creates a message when an incorrect alias is used. The last entry redirects all incoming mail from *olddomain* to *newdomain*.

3. Create the database file.

```
# /usr/sbin/makemap maptype newmap < newmap
```

<i>maptype</i>	Select a database type, such as <i>dbm</i> , <i>btree</i> , or <i>hash</i> .

<i>newmap</i>	Use the name of the input file and the first part of the name of the database file. If the <i>dbm</i> database type is selected, then the database files are created by using a <i>.pag</i> and a <i>.dir</i> suffix. For the other two database types, the file name is followed by <i>.db</i> .

Managing the postmaster Alias

Every system must be able to send mail to a *postmaster* mailbox. You can create an NIS or NIS+ alias for *postmaster*, or you can create the alias in each local */etc/mail/aliases* file. Refer to these procedures.

- “How to Create a postmaster Alias in Each Local `/etc/mail/aliases` File” on page 336
- “How to Create a Separate Mailbox for postmaster” on page 336
- “How to Add the postmaster Mailbox to the Aliases in the `/etc/mail/aliases` File” on page 337

▼ How to Create a postmaster Alias in Each Local `/etc/mail/aliases` File

If you are creating the postmaster alias in each local `/etc/mail/aliases` file, follow these instructions.

1. Become superuser on each local system or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. View the `/etc/mail/aliases` entry.

```
# cat /etc/mail/aliases
# Following alias is required by the mail protocol, RFC 2821
# Set it to the address of a HUMAN who deals with this system's
# mail problems.
Postmaster: root
```

3. Edit each system's `/etc/mail/aliases` file.

Change `root` to the mail address of the person who is designated as the postmaster.

```
Postmaster: mail_address
```

mail_address

Use the assigned address for the person who is designated as the postmaster.

4. (Optional) Create a separate mailbox for the postmaster.

You can create a separate mailbox for the postmaster to keep postmaster mail separate from personal mail. If you create a separate mailbox, use the mailbox address instead of the postmaster's personal mail address when you edit the `/etc/mail/aliases` files. For details, refer to “How to Create a Separate Mailbox for postmaster” on page 336.

▼ How to Create a Separate Mailbox for postmaster

If you are creating a separate mailbox for postmaster, follow these instructions.

1. Become `root` on the mail server or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

Composed September 5, 2002

2. **Create a user account for the person who is designated as `postmaster` and put an asterisk (*) in the password field.**

For details about adding a user account, refer to “Managing User Accounts and Groups (Tasks)” in *System Administration Guide: Basic Administration*.

3. **After mail has been delivered, enable the `mail` program to read and write to the mailbox name.**

```
# mail -f postmaster
```

`postmaster`

Use the assigned address.

▼ How to Add the `postmaster` Mailbox to the Aliases in the `/etc/mail/aliases` File

If you are adding a `postmaster` mailbox to the aliases in the `/etc/mail/aliases` file, follow these instructions.

1. **Become `root` on each system or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Add an alias for `root`. Use the mail address of the person who is designated as the `postmaster`.**

```
# cat /etc/mail/aliases
..
root: user@host.domain.com
```

`user@host.domain.com`

Use the assigned address of the person who is designated as `postmaster`.

3. **On the `postmaster`’s local system, create an entry in the `/etc/mail/aliases` file that defines the name of the alias (`sysadmin`, for example) and include the path to the local mailbox.**

```
# cat /etc/mail/aliases
..
sysadmin: /usr/somewhere/somefile
```

`sysadmin`

Create a name for a new alias.

`/usr/somewhere/somefile`

Use the path to the local mailbox.

4. **Rebuild the alias database.**

newaliases

Administering the Queue Directories (Task Map)

The following table describes the procedures for administering the mail queue.

Task	Description	For Instructions
Displaying the contents of the mail queue, <code>/var/spool/mqueue</code>	Use this procedure to see how many messages are in the queue and how fast the messages are being cleared from the queue.	"How to Display the Contents of the Mail Queue, <code>/var/spool/mqueue</code> " on page 339
Forcing mail queue processing for the mail queue, <code>/var/spool/mqueue</code>	Use this procedure to process messages to a system that previously was unable to receive messages.	"How to Force Mail Queue Processing in the Mail Queue, <code>/var/spool/mqueue</code> " on page 339
Running a subset of the mail queue, <code>/var/spool/mqueue</code>	Use this procedure to force a substring of an address, such as a host name, to be processed or to force a particular message out of the queue.	"How to Run a Subset of the Mail Queue, <code>/var/spool/mqueue</code> " on page 339
Moving the mail queue, <code>/var/spool/mqueue</code>	Use this procedure to move the mail queue.	"How to Move the Mail Queue, <code>/var/spool/mqueue</code> " on page 340
Running the old mail queue, <code>/var/spool/omqueue</code>	Use this procedure to run an old mail queue.	"How to Run the Old Mail Queue, <code>/var/spool/omqueue</code> " on page 341

Administering the Queue Directories (Tasks)

This section describes some helpful tasks for queue administration. For information about the client-only queue, refer to "New Configuration File, `submit.cf`" on page 390. For other related information, you can refer to "New Queue Features" on page 414.

▼ How to Display the Contents of the Mail Queue, /var/spool/mqueue

Use this procedure to see how many messages are in the queue and how fast they are being cleared from the queue.

Use the following command to display this information.

- The queue IDs
- The size of the message
- The date the message entered the queue
- The message status
- The sender and the recipients

```
# /usr/bin/mailq | more
```

This command now checks for the authorization attribute, `solaris.admin.mail.mailq`. If the check is successful, the equivalent of specifying the `-bp` flag with `sendmail` is executed. If the check fails, an error message is printed. By default, this authorization attribute is enabled for all users. The authorization attribute can be disabled by modifying the user entry in `prof_attr`. For more information, refer to the man pages for `prof_attr(4)` and `mailq(1)`.

▼ How to Force Mail Queue Processing in the Mail Queue, /var/spool/mqueue

Use this procedure, for example, to process messages to a system that was previously unable to receive messages.

1. Become root or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Force queue processing and display the progress of the jobs as the queue is cleared.

```
# /usr/lib/sendmail -q -v
```

▼ How to Run a Subset of the Mail Queue, /var/spool/mqueue

Use this procedure, for example, to force a substring of an address, such as a host name, to be processed or to force a particular message from the queue.

1. Become root or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Run a subset of the mail queue at any time with `-qRstring`.

```
# /usr/lib/sendmail -qRstring
```

string

Use a recipient's alias or a substring (like a host name) of *user@host.domain*.

Alternately, you can run a subset of the mail queue with `-qInnnnn`.

```
# /usr/lib/sendmail -qInnnnn
```

nnnnn

Use a queue ID.

▼ How to Move the Mail Queue, `/var/spool/mqueue`

If you are moving the mail queue, follow these instructions.

1. Become `root` on the mail host or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Kill the `sendmail` daemon.

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

Now `sendmail` is no longer processing the queue directory.

3. Change to the `/var/spool` directory.

```
# cd /var/spool
```

4. Move the directory, `mqueue`, and all its contents to the `omqueue` directory. Then create a new empty directory that is named `mqueue`.

```
# mv mqueue omqueue; mkdir mqueue
```

5. Set the permissions of the directory to read/write/execute by owner, and read/execute by group. Also, set the owner and group to `daemon`.

```
# chmod 750 mqueue; chown root:bin mqueue
```

6. Start `sendmail`.

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

▼ How to Run the Old Mail Queue, /var/spool/omqueue

To run an old mail queue, follow these instructions.

1. Become root or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Run the old mail queue.

```
# /usr/lib/sendmail -oQ/var/spool/omqueue -q
```

The `-oQ` flag specifies an alternate queue directory and the `-q` flag says to run every job in the queue. Use the `-v` flag if you are displaying the verbose output on the screen.

3. Remove the empty directory.

```
# rmdir /var/spool/omqueue
```

Administering .forward Files (Task Map)

The following table describes the procedures for administering .forward files. For more information, refer to “.forward Files” on page 379 in Chapter 26.

Task	Description	For Instructions
Disabling .forward files	Use this procedure if, for example, you want to prevent automated forwarding.	“How to Disable .forward Files” on page 342
Changing the .forward file search path	Use this procedure if, for example, you want to move all .forward files into a common directory.	“How to Change the .forward File Search Path” on page 342
Creating and populating /etc/shells	Use this procedure to enable users to use the .forward file to forward mail to a program or to a file.	“How to Create and Populate /etc/shells” on page 343

Administering `.forward` Files (Tasks)

This section contains several procedures that are related to `.forward` file administration. Because these files can be edited by users, they can cause problems. For more information, refer to “`.forward` Files” on page 379 in Chapter 26.

▼ How to Disable `.forward` Files

This procedure, which prevents automated forwarding, disables the `.forward` file for a particular host.

1. Become `root` or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. Make a copy of `/usr/lib/mail/domain/solaris-generic.m4` or your site-specific domain m4 file.

```
# cd /usr/lib/mail/domain
# cp solaris-generic.m4 mydomain.m4
```

mydomain

Use the file name of your choice.

3. Add the following line to the file you just created.

```
define(`confFORWARD_PATH', '') dnl
```

If a value for `confFORWARD_PATH` already exists in the m4 file, replace it with this null value.

4. Build and install a new configuration file.

If you need help with this step, refer to “How to Build a New `sendmail.cf` File” on page 323.

▼ How to Change the `.forward` File Search Path

If, for example, you want to put all `.forward` files in a common directory, follow these instructions.

1. Become `root` or assume an equivalent role.

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

2. **Make a copy of `/usr/lib/mail/domain/solaris-generic.m4` or your site-specific domain m4 file.**

```
# cd /usr/lib/mail/domain
# cp solaris-generic.m4 mydomain.m4
```

mydomain

Use the file name of your choice.

3. **Add the following line to the file that you just created.**

```
define('confFORWARD_PATH', '$z/.forward:/var/forward/$u') dnl
```

If a value for `confFORWARD_PATH` already exists in the m4 file, replace it with this new value.

4. **Build and install a new configuration file.**

If you need help with this step, refer to “How to Build a New `sendmail.cf` File” on page 323.

▼ How to Create and Populate `/etc/shells`

This file is not included in the standard release, so you must add it if users are to be allowed to use `.forward` files to forward mail to a program or to a file. You can create the file manually by using `grep` to identify all of the shells that are listed in your password file. You can then type the shells into the file. However, it is easier to use the following procedure, which employs a script that can be downloaded.

1. **Download the script.**

<http://www.sendmail.org/vendor/sun/gen-etc-shells.html>

2. **Become `root` or assume an equivalent role.**

For information about roles, refer to “Using Privileged Applications” in *System Administration Guide: Security Services*.

3. **To generate a list of shells, run the `gen-etc-shells` script.**

```
# ./gen-etc-shells.sh > /tmp/shells
```

This script uses the `getent` command to collect the names of shells that are included in the password file sources that are listed in `/etc/nsswitch.conf`.

4. **Inspect and edit the list of shells in `/tmp/shells`.**

With the editor of your choice, remove any shells that you are not including.

5. **Move the file to `/etc/shells`.**

```
# mv /tmp/shells /etc/shells
```

Troubleshooting Procedures and Tips for Mail Services (Task Map)

The following table describes troubleshooting procedures and tips for mail services.

Task	Description	For Instructions
Testing mail configuration	Steps for testing changes to the <code>sendmail</code> configuration file	"How to Test the Mail Configuration" on page 344
Checking mail aliases	A step to confirm that mail can or cannot be delivered to a specified recipient	"How to Check Mail Aliases" on page 345
Testing the rule sets	Steps for checking the input and returns of the <code>sendmail</code> rule sets	"How to Test the <code>sendmail</code> Rule Sets" on page 346
Verifying connections to other systems	Tips for verifying connections to other systems	"How to Verify Connections to Other Systems" on page 347
Logging messages by using the <code>syslogd</code> program	Tips for gathering error message information	"Logging Error Messages" on page 347
Checking other sources for diagnostic information	Tips for getting diagnostic information from other sources	"Other Sources for Mail Diagnostic Information" on page 348

Troubleshooting Procedures and Tips for Mail Services (Tasks)

This section provides some procedures and tips that you can use for troubleshooting problems with mail services.

▼ How to Test the Mail Configuration

To test the changes you make to your configuration file, follow these instructions.

1. **Restart `sendmail` on any system that has a revised configuration file.**

```
# pkill -HUP sendmail
```

2. **Send test messages from each system.**


```
# /usr/lib/sendmail -v names </dev/null
```

names

Specify a recipient's email address.

This command sends a null message to the specified recipient and displays the message activity on your monitor.

3. **Send mail to yourself or other people on the local system by addressing the message to a regular user name.**
4. **(Optional) If you are connected to a network, send mail in three directions to someone on another system.**
 - From the main system to a client system
 - From a client system to the main system
 - From a client system to another client system
5. **(Optional) If you have a mail gateway, send mail to another domain from the mail host to ensure that the relay mailer and host are configured properly.**
6. **(Optional) If you have set up a UUCP connection on your phone line to another host, send mail to someone at that host and have that person send mail back or call you when the message is received.**
7. **Ask someone to send mail to you over the UUCP connection.**

The `sendmail` program cannot detect whether the message is delivered because it passes the message to UUCP for delivery.
8. **From different systems, send a message to `postmaster` and ensure that it is delivered to your postmaster's mailbox.**

▼ How to Check Mail Aliases

To verify aliases and to confirm that mail can or cannot be delivered to a specified recipient, follow these instructions.

Display the aliases and identify that the final address is deliverable or not deliverable.

```
% /usr/lib/sendmail -v -bv recipient
```

recipient

Specify a recipient's alias.

The following is an example of the output.

```
% /usr/lib/sendmail -v -bv sandy
sandy... aliased to ssmith
ssmith... aliased to sandy@phoenix
```

```
sandy@phoenix... deliverable: mailer esmtp, host phoenix, user sandy@phoenix.example.com
%
```

Avoid loops and inconsistent databases when both local and domain-wide aliases are used. Be especially careful when you move a user from one system to another to avoid the creation of alias loops.

▼ How to Test the sendmail Rule Sets

To check the input and returns of the sendmail rule sets, follow these instructions.

1. Change to address test mode.

```
# /usr/lib/sendmail -bt
```

2. Test a mail address.

Provide the following numbers and address at the last prompt (>).

```
> 3,0 mail_address
```

mail_address

Use the mail address that you are testing.

3. End the session.

Press Control-d.

The following is an example of the output.

```
% /usr/lib/sendmail -bt
ADDRESS TEST MODE (ruleset 3 NOT automatically invoked)
Enter <ruleset> <address>
> 3,0 sandy@phoenix
canonify          input: sandy @ phoenix
Canonify2         input: sandy < @ phoenix >
Canonify2         returns: sandy < @ phoenix . example . com . >
canonify          returns: sandy < @ phoenix . example . com . >
parse            input: sandy < @ phoenix . example . com . >
Parse0            input: sandy < @ phoenix . example . com . >
Parse0            returns: sandy < @ phoenix . example . com . >
ParseLocal        input: sandy < @ phoenix . example . com . >
ParseLocal        returns: sandy < @ phoenix . example . com . >
Parse1            input: sandy < @ phoenix . example . com . >
MailerToTriple    input: < mailhost . phoenix . example . com >
                  sandy < @ phoenix . example . com . >
MailerToTriple    returns: $# relay @$ mailhost . phoenix . example . com
                  $: sandy < @ phoenix . example . com . >
Parse1            returns: $# relay @$ mailhost . phoenix . example . com
                  $: sandy < @ phoenix . example . com . >
parse            returns: $# relay @$ mailhost . phoenix . example . com
                  $: sandy < @ phoenix . example . com . >
```

▼ How to Verify Connections to Other Systems

The `mconnect` program opens a connection to a mail server on a host that you specify and enables you to test that connection. The program runs interactively, so you can issue various diagnostic commands. See the `mconnect(1)` man page for a complete description. The following example verifies that mail to the user name `sandy` is deliverable.

```
% mconnect phoenix
connecting to host phoenix (172.31.255.255), port 25
connection open
220 phoenix.example.com ESMTP Sendmail 8.12.0+Sun/8.12.0;Sun, 4 Sep 2001 3:52:56 -0700 (PDT)
expn sandy
250 2.1.5 <sandy@phoenix.example.com>
quit
```

If you cannot use `mconnect` to connect to an SMTP port, check these conditions.

- Is the system load too high?
- Is the `sendmail` daemon running?
- Does the system have the appropriate `/etc/mail/sendmail.cf` file?
- Is port 25 (the port that `sendmail` uses) active?

Logging Error Messages

Your mail services logs most error messages by using the `syslogd` program. By default, the `syslogd` program sends these messages to a system that is called `loghost`, which is specified in the `/etc/hosts` file. You can define `loghost` to hold all logs for an entire NIS domain. If no `loghost` is specified, error messages from `syslogd` are not reported.

The `/etc/syslog.conf` file controls where the `syslogd` program forwards messages. You can change the default configuration by editing the `/etc/syslog.conf` file. You must restart the `syslog` daemon for any changes to become active. To gather information about mail, you can add the following selections to the file.

- `mail.alert` – Messages about conditions that should be fixed now
- `mail.crit` – Critical messages
- `mail.warning` – Warning messages
- `mail.notice` – Messages that are not errors, but might need attention
- `mail.info` – Informational messages
- `mail.debug` – Debugging messages

The following entry in the `/etc/syslog.conf` file sends a copy of all critical, informational, and debug messages to `/var/log/syslog`.

```
mail.crit;mail.info;mail.debug          /var/log/syslog
```

Each line in the system log contains a timestamp, the name of the system that generated the line, and a message. The `syslog` file can log a large amount of information.

The log is arranged in a succession of levels. At the lowest level, only unusual occurrences are logged. At the highest level, even the most mundane and uninteresting events are recorded. As a convention, log levels under 10 are considered “useful.” Log levels higher than 10 are usually used for debugging. See the “Customizing System Message Logging” in *System Administration Guide: Advanced Administration* for information about `loghost` and the `syslogd` program.

Other Sources for Mail Diagnostic Information

For other diagnostic information, check the following sources.

- Look at the `Received` lines in the header of the message. These lines trace the route the message took as it was relayed. Remember to consider time–zone differences.
- Look at the messages from `MAILER-DAEMON`. These messages typically report delivery problems.
- Check the system log that records delivery problems for your group of systems. The `sendmail` program always records its activities in the system log. You might want to modify the `crontab` file to run a shell script nightly that searches the log for `SYSERR` messages and mails any messages that it finds to the postmaster.
- Use the `mailstats` program to test mail types and determine the number of incoming and outgoing messages.

Resolving Error Messages

This section describes how you can resolve some `sendmail`–related error messages that are in the Solaris 9 operating environment. You can also refer to <http://www.sendmail.org/faq/>.

The following error messages contain two or more of the following types of information.

- *Cause*: What might have happened to cause the message
- *Description*: What the user was doing when the error message occurred
- *Action*: What you can do to fix the problem or to continue with your work
- *Technical Notes*: Background information that might be interesting or helpful to a technical audience, such as developers

- *See Also:* Suggests further reading

451 timeout waiting for input during *source*

Cause: When sendmail reads from any source that might time out, such as an SMTP connection, it sets a timer to the value of various `Timeout` options before reading begins. If the read is not completed before the timer expires, this message appears and reading stops. (Usually this situation occurs during `RCPT`.) The mail message is then queued for later delivery.

Action: If you see this message often, increase the value of various `Timeout` options in the `/etc/mail/sendmail.cf` file. If the timer is already set to a large number, look for hardware problems, such as poor network cabling or connections.

See Also: For more information about the `Timeout` option, refer to “Changes to the `Timeout` Option” on page 404. If you are using online documentation, the term “timeouts” is a good search string.

550 *hostname*... Host unknown

Cause: This sendmail message indicates that the destination host machine, which is specified by the portion of the address after the at sign (`@`), was not found during domain name system (DNS) lookup.

Action: Use the `nslookup` command to verify that the destination host exists in that domain or other domains, perhaps with a slightly different spelling. Otherwise, contact the intended recipient and ask for a proper address.

550 *username*... User unknown

Cause: This sendmail message indicates that the intended recipient, who is specified by the portion of the address before the at sign (`@`), could not be located on the destination host machine.

Action: Check the email address and try again, perhaps with a slightly different spelling. If this remedy does not work, contact the intended recipient and ask for a proper address.

554 *hostname*... Local configuration error

Cause: This sendmail message usually indicates that the local host is trying to send mail to itself.

Action: Check the value of the `$j` macro in the `/etc/mail/sendmail.cf` file to ensure that this value is a fully qualified domain name.

Technical Notes: When the sending system provides its host name to the receiving system (in the SMTP `HELO` command), the receiving system compares its name to the sender’s name. If these names are the same, the receiving system issues this error message and closes the connection. The name that is provided in the `HELO` command is the value of the `$j` macro.

Composed September 5, 2002

See Also: For additional information, refer to <http://www.sendmail.org/faq/section4.html#4.5>.

config error: mail loops back to myself.

Cause: If you set up an MX record and make host *bar* the mail exchanger for domain *foo*, but you do not configure host *bar* to know that it is the mail exchanger for domain *foo*, you receive this error message.

Also, another possibility is that both the sending system and the receiving system are identifying as the same domain.

Action: For instructions, refer to <http://www.sendmail.org/faq/section4.html#4.5>.

host name configuration error

Action: Follow the instructions that were provided for resolving this error message, 554 *hostname...* Local configuration error.

Technical Notes: This is an old sendmail message, which replaced I refuse to talk to myself and is now replaced by the Local configuration error message.

user unknown

Description: When you try to send mail to a user, the error Username... user unknown is displayed. The user is on the same system.

Action: Check for a typographical error in the entered email address. Otherwise, the user could be aliased to a nonexistent email address in */etc/mail/aliases* or in the user's *.mailrc* file. Also, check for uppercase characters in the user name. Preferably, email addresses should not be case sensitive.

See Also: For additional information, refer to <http://www.sendmail.org/faq/section4.html#4.17>.

Mail Services (Reference)

The `sendmail` program is a mail transport agent that uses a configuration file to provide aliasing and forwarding, automatic routing to network gateways, and flexible configuration. The Solaris operating environment supplies standard configuration files that most sites can use. Chapter 24 provided an introduction to the components of mail services and a description of a typical mail service configuration. Chapter 25 explained how to set up and administer an electronic mail system. This chapter provides information on the following topics.

- “Solaris Version of `sendmail`” on page 351
- “Software and Hardware Components of Mail Services” on page 355
- “Mail Service Programs and Files” on page 365
- “Mail Addresses and Mail Routing” on page 382
- “Interactions of `sendmail` With Name Services” on page 383

See Chapter 27 for a description of the new features that are included in version 8.12 of `sendmail`, the version that is in this Solaris 9 release. You can also read about changes to `mail.local`, `mailstats`, `makemap`, and about a new maintenance utility, `editmap`. For details that are not covered in these chapters, see the man pages for `sendmail(1M)`, `mail.local(1M)`, `mailstats(1)`, `makemap(1M)`, and `editmap(1M)`.

Solaris Version of `sendmail`

This section, which includes the following topics, describes some of the differences in the Solaris version of `sendmail` as compared to the generic Berkeley version.

- “Flags Used and Not Used to Compile `sendmail`” on page 352
- “Alternative `sendmail` Commands” on page 353
- “Versions of the Configuration File” on page 354

Flags Used and Not Used to Compile `sendmail`

The following tables list the flags that are used when compiling the version of `sendmail` that is delivered with the Solaris 9 release. If your configuration requires other flags, you need to download the source and recompile the binary. You can find information about this process at <http://www.sendmail.org>.

TABLE 26-1 General `sendmail` Flags

Flag	Description
<code>SOLARIS=20900</code>	Support for the Solaris 9 operating environment.
<code>MILTER</code>	Support for the Mail Filter API.
<code>NETINET6</code>	Support for IPv6. This flag has been moved from <code>conf.h</code> to <code>Makefile</code> .

TABLE 26-2 Maps and Database Types

Flag	Description
<code>NDBM</code>	Support for <code>ndbm</code> databases
<code>NEWDB</code>	Support for <code>db</code> databases
<code>USERDB</code>	Support for the User database
<code>NIS</code>	Support for <code>nis</code> databases
<code>NISPLUS</code>	Support for <code>nisplus</code> databases
<code>LDAPMAP</code>	Support for LDAP maps
<code>MAP_REGEX</code>	Support for regular expression maps

TABLE 26-3 Solaris Flags

Flag	Description
<code>SUN_EXTENSIONS</code>	Support for Sun extensions that are included in <code>sun_compat.o</code> .
<code>SUN_LOOKUP_MACRO</code>	Support for the L and G configuration commands in <code>sendmail.cf</code> . Use of these commands is not recommended.
<code>SUN_DEFAULT_VALUES</code>	Support only for the default values in the Solaris flag, <code>SUN_CONTENT_LENGTH</code> .

TABLE 26-3 Solaris Flags (Continued)

Flag	Description
SUN_INIT_DOMAIN	For backward compatibility, support for the use of NIS domain names to fully qualify the local host name. For more information, look for vendor-specific information in http://www.sendmail.org .
SUN_CONTENT_LENGTH	Support for the Content-Length: header in messages to files. For more information, look for vendor-specific information in http://www.sendmail.org .
SUN_SIMPLIFIED_LDAP	Support for a simplified LDAP API, which is specific to Sun. For more information, look for vendor-specific information in http://www.sendmail.org .
VENDOR_DEFAULT=VENDOR_SUN	Selects Sun as the default vendor.

The following table lists generic flags that are not used to compile the version of `sendmail` that is delivered with the Solaris 9 release.

TABLE 26-4 Generic Flags Not Used in the Solaris Version of `sendmail`

Flag	Description
SASL	Simple Authentication and Security Layer (RFC 2554)
STARTTLS	Transaction Level Security (RFC 2487)

To see a list of the flags that are used to compile `sendmail`, use the following command.

```
% /usr/lib/sendmail -bt -d0.10 < /dev/null
```

Note – The preceding command does not list the flags that are specific to Sun.

Alternative `sendmail` Commands

The Solaris release does not include all of the command synonyms that are provided in the generic release from Berkeley. This table includes a complete list of the command aliases, whether they are included in the Solaris release, and how to generate the same behavior by using `sendmail`.

TABLE 26-5 Alternate `sendmail` Commands

Alternate Name	In the Solaris Release?	Options With <code>sendmail</code>
<code>hoststat</code>	No	<code>sendmail -bh</code>
<code>mailq</code>	Yes	<code>sendmail -bp</code>
<code>newaliases</code>	Yes	<code>sendmail -bi</code>
<code>purgestat</code>	No	<code>sendmail -bH</code>
<code>smtpd</code>	No	<code>sendmail -bd</code>

Versions of the Configuration File

The Solaris 9 version of `sendmail` includes a configuration option that enables you to define the version of the `sendmail.cf` file. This option enables older configuration files to be used with the current version of `sendmail`. You can set the version level to values between 0 and 10. You can also define the vendor. Either Berkeley or Sun are valid vendor options. If a version level is specified but no vendor is defined, Sun is used as the default vendor setting. The following table lists some of the valid options.

TABLE 26-6 Configuration File Version Values

Field	Description
V7/Sun	Setting that was used for Version 8.8 of <code>sendmail</code> .
V8/Sun	Setting that was used for Version 8.9 of <code>sendmail</code> . This setting was included in the Solaris 8 release.
V9/Sun	Setting that was used for versions 8.10 and 8.11 of <code>sendmail</code> .
V10/Sun	Setting that is used for version 8.12 of <code>sendmail</code> . Version 8.12 is the default for the Solaris 9 release.

Note – You are urged not to use V1/Sun. For more information, refer to <http://www.sendmail.org/vendor/sun/differences.html#4>.

For task information, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

Software and Hardware Components of Mail Services

This section describes the software and hardware components of a mail system.

- “Software Components” on page 355
- “Hardware Components” on page 362

Software Components

Each mail service includes at least one of each of the following software components.

- “Mail User Agent” on page 355
- “Mail Transfer Agent” on page 355
- “Local Delivery Agent” on page 356

This section also describes these software components.

- “Mailers” on page 356
- “Mail Addresses” on page 357
- “Mailbox Files” on page 360
- “Mail Aliases” on page 361

Mail User Agent

The *mail user agent* is the program that acts as the interface between the user and mail transfer agent. The `sendmail` program is a mail transfer agent. The Solaris operating environment supplies the following mail user agents.

- `/usr/bin/mail`
- `/usr/bin/mailx`
- `$OPENWINHOME/bin/mailtool`
- `/usr/dt/bin/dtmail`

Mail Transfer Agent

The *mail transfer agent* is responsible for the routing of mail messages and the resolution of mail addresses. This agent is also known as a mail *transport* agent. The transfer agent for the Solaris operating environment is `sendmail`. The transfer agent performs these functions.

- Accepts messages from the mail user agent

- Resolves destination addresses
- Selects a proper delivery agent to deliver the mail
- Receives incoming mail from other mail transfer agents

Local Delivery Agent

A *local delivery agent* is a program that implements a mail delivery protocol. The following local delivery agents are provided with the Solaris operating environment.

- The UUCP local delivery agent, which uses `uux` to deliver mail
- The local delivery agent, which is `mail.local` in the standard Solaris release

Chapter 27 provides information on these related topics.

- “New Delivery Agent Flags” on page 412
- “New Equates for Delivery Agents” on page 413
- “Changes to `mail.local`” on page 418

Mailers

Mailer is a `sendmail`-specific term. A *mailer* is used by `sendmail` to identify a specific instance of a customized local delivery agent or a customized mail transfer agent. You need to specify at least one mailer in your `sendmail.cf` file. For task information, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25. This section provides a brief description of two types of mailers.

- “Simple Mail Transfer Protocol (SMTP) Mailers” on page 356
- “UNIX-to-UNIX Copy Program (UUCP) Mailers” on page 357

For additional information about mailers, see
<http://www.sendmail.org/m4/readme.html> or `/usr/lib/mail/README`.

Simple Mail Transfer Protocol (SMTP) Mailers

SMTP is the standard mail protocol that is used on the Internet. This protocol defines these mailers.

- `smtp` provides regular (historic-style) SMTP transfers to other servers
- `esmtplib` provides extended SMTP transfers to other servers
- `smtp8` provides SMTP transfers to other servers without converting 8-bit data to MIME
- `dsmtplib` provides on-demand delivery by using the `F=%` mailer flag. Refer to “Changes to the `MAILER()` Declaration” on page 412 and “New Delivery Agent Flags” on page 412 in Chapter 27.

UNIX-to-UNIX Copy Program (UUCP) Mailers

If possible, avoid using UUCP. For an explanation, refer to <http://www.sendmail.org/m4/uucp.html> or do a search in `/usr/lib/mail/README` on this string: `USING UUCP MAILERS`.

UUCP defines these mailers.

<code>uucp-old</code>	Names in the <code>\$=U</code> class are sent to <code>uucp-old</code> . <code>uucp</code> is the obsolete name for this mailer. The <code>uucp-old</code> mailer uses an exclamation-point address in the headers.
<code>uucp-new</code>	Names in the <code>\$=Y</code> class are sent to <code>uucp-new</code> . Use this mailer when you know that the receiving UUCP mailer can manage multiple recipients in one transfer. <code>suucp</code> is the obsolete name for this mailer. The <code>uucp-new</code> mailer also uses an exclamation-point address in the headers.

If `MAILER (smtp)` is also specified in your configuration, two more mailers are defined.

<code>uucp-dom</code>	This mailer uses domain-style addresses and, basically, applies the SMTP rewriting rules.
<code>uucp-uudom</code>	Names in the <code>\$=Z</code> class are sent to <code>uucp-uudom</code> . <code>uucp-uudom</code> and <code>uucp-dom</code> use the same header address format, domain-style addresses.

Note – Because the `smtp` mailer modifies the UUCP mailer, always put `MAILER (smtp)` before `MAILER (uucp)` in your `.mc` file.

Mail Addresses

The *mail address* contains the name of the recipient and the system to which the mail message is delivered. When you administer a small mail system that does not use a name service, addressing mail is easy. The login names uniquely identify the users. Complexity is introduced if you are administering a mail system that has more than one system with mailboxes or has one or more domains. Additional complexity can be generated if you have a UUCP (or other) mail connection to the outside world. The information in the following sections can help you understand the parts and complexities of a mail address.

- “Domains and Subdomains” on page 358
- “Name Service Domain Name and Mail Domain Name” on page 358
- “Typical Format for Mail Addresses” on page 359
- “Route-Independent Mail Addresses” on page 359

Domains and Subdomains

Email addressing uses domains. A *domain* is a directory structure for network address naming. A domain can have one or more *subdomains*. The domain and subdomains of an address can be compared to the hierarchy of a file system. Just as a subdirectory is considered to be inside the directory above it, each subdomain in a mail address is considered to be inside the location to its right.

The following table shows some top-level domains.

TABLE 26-7 Top-Level Domains

Domain	Description
com	Commercial sites
edu	Educational sites
gov	United States government installations
mil	United States military installations
net	Networking organizations
org	Other nonprofit organizations

Domains are case insensitive. You can use uppercase, lowercase, or mixed-case letters in the domain part of an address without making any difference.

For more information about domains, refer to “Domain Name System (Overview)” in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Name Service Domain Name and Mail Domain Name

When you are working with name service domain names and mail domain names, remember the following.

- By default, the `sendmail` program strips the first component from the NIS or NIS+ domain name to form the mail domain name. For example, if an NIS+ domain name were `bldg5.example.com`, its mail domain name would be `example.com`.
- Although mail domain addresses are case insensitive, the NIS or NIS+ domain name is not. For the best results, use lowercase characters when setting up the mail and NIS or NIS+ domain names.
- The DNS domain name and the mail domain name must be identical.

For more information, refer to “Interactions of `sendmail` With Name Services” on page 383.

Typical Format for Mail Addresses

Typically, a mail address has the following format. For further details, refer to “Route–Independent Mail Addresses” on page 359.

user@subdomain. subdomain2.subdomain1.top-level-domain

The part of the address to the left of the @ sign is the local address. The local address can contain the following.

- Information about routing with another mail transport (for example, bob: :vmsvax@gateway or smallberries%mill.uucp@gateway)
- An alias (for example, iggy.ignatz)

Note – The receiving mailer is responsible for determining what the local part of the address means. For information about mailers, refer to “Mailers” on page 356.

The part of the address to the right of the @ sign shows the domain levels, which is where the local address resides. A dot separates each subdomain. The domain part of the address can be an organization, a physical area, or a geographic region. Furthermore, the order of domain information is hierarchical—the more local the subdomain, the closer it is to the @ sign.

Route–Independent Mail Addresses

Mail addresses can be route independent. Route-independent addressing requires the sender of an email message to specify the name of the recipient and the final destination. A high-speed network, such as the Internet, uses route-independent addresses. Route-independent addresses can have this format.

user@host . domain

Route-independent addresses for UUCP connections can have this address format.

host . domain ! user

The increased popularity of the domain-hierarchical naming scheme for computers is making route-independent addresses more common. Actually, the most common route-independent address omits the host name and relies on the domain name service to properly identify the final destination of the email message.

user@domain

Route-independent addresses are read by searching for the @ sign and then reading the domain hierarchy from the right (the highest level) to the left (the most specific part of the address to the right of the @ sign).

Mailbox Files

A *mailbox* is a file that is the final destination for email messages. The name of the mailbox can be the user name or the identity of a specific function, such as the postmaster. Mailboxes are in the `/var/mail/username` file, which can exist either on the user's local system or on a remote mail server. In either instance, the mailbox is on the system to which the mail is delivered.

Mail should always be delivered to a local file system so that the user agent can pull mail from the mail spool and store it readily in the local mailbox. Do not use NFS-mounted file systems as the destination for a user's mailbox. Specifically, do not direct mail to a mail client that is mounting the `/var/mail` file system from a remote server. Mail for the user, in this instance, should be addressed to the mail server and not to the client host name. NFS-mounted file systems can cause problems with mail delivery and handling.

The `/etc/mail/aliases` file and name services such as NIS and NIS+ provide mechanisms for creating aliases for electronic mail addresses so that users do not need to know the precise local name of a user's mailbox.

The following table shows some common naming conventions for special-purpose mailboxes.

TABLE 26-8 Conventions for the Format of Mailbox Names

Format	Description
<i>username</i>	User names are frequently the same as mailbox names.
<i>Firstname.Lastname</i> <i>Firstname_Lastname</i> <i>Firstinitial.Lastname</i> <i>Firstinitial_Lastname</i>	User names can be identified as full names with a dot (or an underscore) that separates the first and last names, or by a first initial with a dot (or an underscore) that separates the initial and the last name.
<i>postmaster</i>	Users can address questions and report problems with the mail system to the <i>postmaster</i> mailbox. Each site and domain should have a <i>postmaster</i> mailbox.
<i>MAILER-DAEMON</i>	<i>sendmail</i> automatically routes any mail that is addressed to the <i>MAILER-DAEMON</i> to the <i>postmaster</i> .
<i>aliasname-request</i>	Names that end in <i>-request</i> are administrative addresses for distribution lists. This address should redirect mail to the person who maintains the distribution list.
<i>owner-aliasname</i>	Names that begin with <i>owner-</i> are administrative addresses for distribution lists. This address should redirect mail to the person who handles mail errors.
<i>owner-owner</i>	This alias is used when no <i>owner-aliasname</i> alias exists for errors to be returned to. This address should redirect mail to the person who handles mail errors and should be defined on any system that maintains a large number of aliases.

TABLE 26-8 Conventions for the Format of Mailbox Names (Continued)

Format	Description
<i>local%domain</i>	The percent sign (%) marks a local address that is expanded when the message arrives at its destination. Most mail systems interpret mailbox names with % characters as full mail addresses. The % is replaced with an @, and the mail is redirected accordingly. Although many people use the % convention, it is not a formal standard. This convention is referred to as the “percent hack.” This feature is often used to help debug mail problems.

Starting with `sendmail` version 8, the envelope sender for mail that is sent to a group alias has been changed to the address that is expanded from the owner alias, if an owner alias exists. This change enables any mail errors to be sent to the alias owner, rather than being returned to the sender. With this change, users notice that mail that was sent to an alias looks as if it came from the alias owner, when delivered. The following alias format helps with some of the problems that are associated with this change.

```
mygroup: :include:/pathname/mygroup.list
owner-mygroup: mygroup-request
mygroup-request: sandys, ignatz
```

In this example, the `mygroup` alias is the actual mail alias for the group. The `owner-mygroup` alias receives error messages. The `mygroup-request` alias should be used for administrative requests. This structure means that in mail sent to the `mygroup` alias, the envelope sender changes to `mygroup-request`.

Mail Aliases

An *alias* is an alternate name. For email, you can use aliases to assign a mailbox location or to define mailing lists. For a task map, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25. Also, you can refer to “Mail Alias Files” on page 376 in this chapter.

For large sites, the mail alias typically defines the location of a mailbox. Providing a mail alias is like providing a room number as part of the address for an individual at a large corporation that occupies multiple rooms. If you do not provide the room number, the mail is delivered to a central address. Without a room number, extra effort is required to determine where within the building the mail is to be delivered, and the possibility of an error increases. For example, if two people who are named Kevin Smith are in the same building, only one of them might get mail. To correct the problem, each Kevin Smith should have a room number added to his address.

Use domains and location-independent addresses as much as possible when you create mailing lists. To enhance portability and flexibility of alias files, make your alias entries in mailing lists as generic and system independent as possible. For example, if you have a user who is named `ignatz` on system `mars`, in domain `example.com`,

create the alias `ignatz@example` instead of `ignatz@mars`. If user `ignatz` changes the name of his system but remains within the `example` domain, you do not need to update alias files to reflect the change in system name.

When you create alias entries, type one alias per line. You should have only one entry that contains the user's system name. For example, you could create the following entries for user `ignatz`.

```
ignatz: iggy.ignatz
iggyi: iggy.ignatz
iggy.ignatz: ignatz@mars
```

You can create an alias for local names or domains. For example, an alias entry for user `fred`, who has a mailbox on the system `mars` and is in the domain `planets`, could have this entry in the NIS+ aliases table.

```
fred: fred@planets
```

When you create mail lists that include users outside your domain, create the alias with the user name and the domain name. For example, if you have a user named `smallberries` on system `privet`, in domain `example.com`, create the alias as `smallberries@example.com`. The email address of the sender is now automatically translated to a fully qualified domain name when mail goes outside the user's domain.

The following list describes methods for creating and administering mail alias files.

- You can create mail aliases for global use in the NIS+ `mail_aliases` table, the NIS `aliases` map, or in local `/etc/mail/aliases` files. You can also create and administer mailing lists that use the same alias files.
- Depending on the configuration of your mail services, you can administer aliases by using the NIS or NIS+ name service to maintain a global `aliases` database or by updating all the local `/etc/mail/aliases` files to keep them synchronized.
- Users can also create and use aliases. They can create aliases either in their local `~/.mailrc` file, which only they can use, or in their local `/etc/mail/aliases` file, which anyone can use. Users cannot normally create or administer NIS or NIS+ alias files.

Hardware Components

You can provide the three required elements of mail configuration in the same system or have separate systems provide these elements.

- "Mail Host" on page 363
- "Mail Server" on page 363
- "Mail Client" on page 364

When users are to communicate with networks outside your domain, you must also add a fourth element, a mail gateway. For more information, refer to “Mail Gateway” on page 364. The following sections describe each hardware component.

Mail Host

A *mail host* is the machine that you designate as the main mail machine on your network. A mail host is the machine to which other systems at the site forward mail that they cannot deliver. You designate a system as a mail host in the `hosts` database by adding the word `mailhost` to the right of the IP address in the local `/etc/hosts` file or in the `hosts` file in the name service. You must also use the `main.cf` file as the mail configuration file on the mail host system. For detailed task information, refer to “How to Set Up a Mail Host” on page 319 in Chapter 25.

A good candidate for a mail host is a system on the local area network that also has a modem for setting up PPP or UUCP links over telephone lines. Another good candidate is a system that is configured as a router from your network to the Internet global network. For more information, refer to Chapter 29, Chapter 38, and “Configuring Routers” in *System Administration Guide: IP Services*. If none of the systems on your local network has a modem, designate one as the mail host.

Some sites use standalone machines that are not networked in a time-sharing configuration. That is, the standalone machine serves terminals that are attached to its serial ports. You can set up electronic mail for this configuration by designating the standalone system as the mail host of a one-system network. “Overview of the Hardware Components” on page 307 in Chapter 24 provides a figure that shows a typical email configuration.

Mail Server

A *mailbox* is a single file that contains email for a particular user. Mail is delivered to the system where the user’s mailbox resides, which can be on a local machine or a remote server. A *mail server* is any system that maintains user mailboxes in its `/var/mail` directory. For task information, refer to “How to Set Up a Mail Server” on page 315 in Chapter 25.

The mail server routes all mail from a client. When a client sends mail, the mail server puts it in a queue for delivery. After the mail is in the queue, a user can reboot or turn off the client without losing those mail messages. When the recipient gets mail from a client, the path in the “From” line of the message contains the name of the mail server. If the recipient responds, the response goes to the user’s mailbox. Good candidates for mail servers are systems that provide a home directory for users or systems that are backed up regularly.

If the mail server is not the user’s local system, users in configurations that use NFS software can mount the `/var/mail` directory by using the `/etc/vfstab` file (if they have root access) or by using the automounter. If NFS support is not available, users can log in to the server to read their mail.

If users on your network send other types of mail, such as audio files or files from desktop publishing systems, you need to allocate more space on the mail server for mailboxes.

One advantage to establishing a mail server for all mailboxes is that it simplifies backups. Backups can be difficult to do when mail is spread over many systems. The disadvantage of storing many mailboxes on one server is that the server can be a single point of failure for many users. However, the advantages of providing good backups usually make the risk worthwhile.

Mail Client

A *mail client* is any system that receives mail on a mail server and does not have a local `/var/mail` directory. This type of configuration is known as remote mode. Remote mode is enabled by default in `/etc/mail/subsidiary.cf`.

You must check that the mail client has the appropriate entry in the `/etc/vfstab` file and a mount point to mount the mailbox from the mail server. Also, ensure that the alias for the client is directed to the mail server's host name, not to the client's name. For task information, refer to "How to Set Up a Mail Client" on page 317 in Chapter 25.

Mail Gateway

The *mail gateway* is a machine that handles connections between networks that run different communications protocols or communications between different networks that use the same protocol. For example, a mail gateway might connect a TCP/IP network to a network that runs the Systems Network Architecture (SNA) protocol suite.

The simplest mail gateway to set up is one that connects two networks that use the same protocol or mailer. This system handles mail with an address for which `sendmail` cannot find a recipient in your domain. If a mail gateway exists, `sendmail` uses it for sending and receiving mail outside your domain.

You can set up a mail gateway between two networks that use unmatched mailers, as shown in the next figure. To support this configuration, you must customize the `sendmail.cf` file on the mail gateway system, which can be a difficult and time-consuming process.




FIGURE 26-1 Gateway Between Different Communications Protocols

If you have to set up a mail gateway, you should find a gateway configuration file that is close to what you need and modify it to conform to your situation.

If you have a machine that provides connections to the Internet, you can configure that machine as the mail gateway. Carefully consider your site's security needs before you configure a mail gateway. You might need to create a firewall gateway between your corporate network and the outside world, and set up that gateway as the mail gateway. For task information, refer to "How to Set Up a Mail Gateway" on page 320 in Chapter 25.

Mail Service Programs and Files

Mail services include many programs and daemons that interact with each other. This section introduces the files, programs, terms, and concepts that are related to administering electronic mail.

- "Contents of the `/usr/bin` Directory" on page 366
- "Contents of the `/etc/mail` Directory" on page 366
- "Contents of the `/usr/lib` Directory" on page 367
- "Other Files Used for Mail Services" on page 370
- "Interactions of Mail Programs" on page 370
- "sendmail Program" on page 371
- "Mail Alias Files" on page 376
- ".forward Files" on page 379
- "`/etc/default/sendmail` File" on page 381

Contents of the `/usr/bin` Directory

The following table shows the contents of the `/usr/bin` directory, which is used for mail services.

Name	Type	Description
<code>aliasadm</code>	File	A program to manipulate the NIS+ aliases map.
<code>mail</code>	File	A user agent.
<code>mailcompat</code>	File	A filter to store mail in SunOS 4.1 mailbox format.
<code>mailq</code>	Link	A link to <code>/usr/lib/sendmail</code> . Used to list the mail queue.
<code>mailstats</code>	File	A program that is used to read mail statistics that are stored in the <code>/etc/mail/sendmail.st</code> file (if present).
<code>mailx</code>	File	A user agent.
<code>mconnect</code>	File	A program that connects to the mailer for address verification and debugging.
<code>praliases</code>	File	A command to “uncompile” the alias database. Refer to the uncompile information that is provided in the man page for <code>praliases(1)</code> .
<code>rmail</code>	Link	A link to <code>/usr/bin/mail</code> . Command that is often used to permit only the sending of mail.
<code>vacation</code>	File	A command to set up an automatic reply to mail.

Contents of the `/etc/mail` Directory

The following table shows the contents of the `/etc/mail` directory.

Name	Type	Description
<code>Mail.rc</code>	File	Default settings for the <code>mailtool</code> user agent.
<code>aliases</code>	File	Mail-forwarding information.
<code>aliases.db</code>	File	Default binary form of mail-forwarding information (created by running <code>newaliases</code>).
<code>aliases.dir</code>	File	Binary form of mail-forwarding information (created by running <code>newaliases</code>). Can still be used, but is no longer used by default in the Solaris 9 release.

Name	Type	Description
aliases.pag	File	Binary form of mail-forwarding information (created by running <code>newaliases</code>). Can still be used, but is no longer used by default in the Solaris 9 release.
mailx.rc	File	Default settings for the <code>mailx</code> user agent.
main.cf	File	Sample configuration file for main systems.
relay-domains	File	List of all domains for which relaying is allowed. By default, only the local domain is allowed.
sendmail.cf	File	Configuration file for mail routing.
submit.cf	File	New configuration file for the mail submission program (MSP). For more information, refer to “New Configuration File, <code>submit.cf</code> ” on page 390.
local-host-names	File	Optional file that you can create if the number of aliases for the mail host is too long.
helpfile	File	Help file that is used by the SMTP <code>HELP</code> command.
sendmail.pid	File	File that lists the PID of the listening daemon and is now in <code>/var/run</code> .
sendmail.st	File	<code>sendmail</code> statistics file. If this file is present, <code>sendmail</code> logs the amount of traffic through each mailer.
subsidiary.cf	File	Sample configuration file for subsidiary systems.
trusted-users	File	File that lists the users (one per line) who can be trusted to perform certain mail operations. By default, only <code>root</code> is in this file. Certain mail operations, when performed by untrusted users, result in the following warning, <code>X-Authentication-Warning:</code> header being added to a message.

Contents of the `/usr/lib` Directory

Table 26–9 shows the contents of the `/usr/lib` directory, which is used for mail services.

TABLE 26–9 Contents of the `/usr/lib` Directory

Name	Type	Description
mail.local	File	Mailer that delivers mail to mailboxes.
sendmail	File	Routing program, also known as the mail transfer agent.

TABLE 26-9 Contents of the `/usr/lib` Directory (Continued)

Name	Type	Description
smrsh	File	Shell program (sendmail restricted shell) that uses the " <code> program</code> " syntax of <code>sendmail</code> to restrict programs that <code>sendmail</code> can run to those programs listed in the <code>/var/adm/sm.bin</code> directory. Refer to the <code>smrsh(1M)</code> man page for recommendations on what to include in <code>/var/adm/sm.bin</code> . To enable, include this <code>m4</code> command, <code>FEATURE('smrsh')</code> , in your <code>mc</code> file.

Contents of the `/usr/lib/mail` Directory

Within the `/usr/lib` directory is a subdirectory, `mail`, that contains all of the necessary files to build a `sendmail.cf` file. The contents of `mail` are shown in Table 26-10.

TABLE 26-10 Contents of the `/usr/lib/mail` Directory Used for Mail Services

Name	Type	Description
README	File	Describes the configuration files.
cf	Directory	Provides site-dependent and site-independent descriptions of hosts.
cf/main.mc	File	Previously named <code>cf/main-v7sun.mc</code> . Is the main configuration file.
cf/makefile	File	Provides rules for building new configuration files.
cf/submit.mc	File	Is the configuration file for the mail submission program (MSP), which is used to submit messages.
cf/subsidiary.mc	File	Previously named <code>cf/subsidiary-v7sun.mc</code> . Is the configuration file for hosts that NFS-mount <code>/var/mail</code> from another host.
domain	Directory	Provides site-dependent subdomain descriptions.
domain/generic.m4	File	Is the generic domain file from Berkeley.

TABLE 26-10 Contents of the `/usr/lib/mail` Directory Used for Mail Services
(Continued)

Name	Type	Description
<code>domain/solaris-antispam.m4</code>	File	Is the domain file with changes that make <code>sendmail</code> function like previous Solaris versions. However, relaying is disabled completely, sender addresses with no host name are rejected, and unresolvable domains are rejected.
<code>domain/solaris-generic.m4</code>	File	Is the default domain file with changes that make <code>sendmail</code> function like previous Solaris versions.
<code>feature</code>	Directory	Contains definitions of specific features for particular hosts (see <code>README</code> for a full description of the features).
<code>m4</code>	Directory	Contains site-independent include files.
<code>mailer</code>	Directory	Contains definitions of mailers, which include <code>local</code> , <code>smtp</code> , and <code>uucp</code> .
<code>ostype</code>	Directory	Describes various operating system environments.
<code>ostype/solaris2.m4</code>	File	Defines default local mailer as <code>mail.local</code> .
<code>ostype/solaris2.ml.m4</code>	File	Defines default local mailer as <code>mail.local</code> .
<code>ostype/solaris2.pre5.m4</code>	File	Defines local mailer as <code>mail</code> .
<code>ostype/solaris8.m4</code>	File	Defines local mailer as <code>mail.local</code> (in LMTP mode), enables IPv6, specifies <code>/var/run</code> as the directory for the <code>sendmail.pid</code> file.
<code>sh</code>	Directory	Contains shell scripts that are used by the <code>m4</code> build process and migration aids.
<code>sh/check-permissions</code>	File	Checks permissions of <code>:include:</code> aliases and <code>.forward</code> files and their parent directory path for correct permissions.
<code>sh/check-hostname</code>	File	Verifies that <code>sendmail</code> is able to determine the fully qualified host name.

Other Files Used for Mail Services

Several other files and directories are used for mail services, as shown in Table 26–11.

TABLE 26–11 Other Files Used for Mail Services

Name	Type	Description
<code>sendmailvars.org_dir</code>	Table	NIS+ version of <code>sendmailvars</code> file.
<code>/etc/default/sendmail</code>	File	Lists the environment variables for the startup script for <code>sendmail</code> .
<code>/etc/shells</code>	File	Lists the valid login shells.
<code>/usr/sbin/editmap</code>	File	Queries and edits single records in database maps for <code>sendmail</code> .
<code>/usr/sbin/in.comsat</code>	File	Mail notification daemon.
<code>/usr/sbin/makemap</code>	File	Builds binary forms of keyed maps.
<code>/usr/sbin/newaliases</code>	Link	A link to <code>/usr/lib/sendmail</code> . Used to create the binary form of the alias database. Previously in <code>/usr/bin</code> .
<code>/usr/sbin/syslogd</code>	File	Error message logger, used by <code>sendmail</code> .
<code>/usr/sbin/etrn</code>	File	Perl script for starting the client-side remote mail queue.
<code>/usr/dt/bin/dtmail</code>	File	CDE mail user agent.
<code>/var/mail/mailbox1</code> , <code>/var/mail/mailbox2</code>	File	Mailboxes for delivered mail.
<code>/var/spool/clientmqueue</code>	Directory	Storage for mail that is delivered by the client daemon.
<code>/var/spool/mqueue</code>	Directory	Storage for mail that is delivered by the master daemon.
<code>\$OPENWINHOME/bin/mailtool</code>	File	Window-based mail user agent.
<code>/var/run/sendmail.pid</code>	File	File that lists the PID of the listening daemon.

Interactions of Mail Programs

Mail services are provided by a combination of the following programs, which interact as shown in the simplified illustration in Figure 26–2.




FIGURE 26-2 Interactions of Mail Programs

For a more detailed illustration, refer to the figure in “sendmail Features” on page 374.

The following is a description of the interactions of mail programs.

1. Users send messages by using programs such as `mailx` or `mailtool`. See the man pages for `mailx(1)` or `mailtool(1)` for information about these programs.
2. The message is collected by the program that generated it and is passed to the `sendmail` daemon.
3. The `sendmail` daemon *parses* the addresses (divides them into identifiable segments) in the message. The daemon uses information from the configuration file, `/etc/mail/sendmail.cf`, to determine network name syntax, aliases, forwarding information, and network topology. By using this information, `sendmail` determines the route a message must follow to get to a recipient.
4. The `sendmail` daemon passes the message to the appropriate system.
5. The `/usr/lib/mail.local` program on the local system delivers the mail to the mailbox in the `/var/mail/username` directory of the recipient of the message.
6. The recipient is notified that mail has arrived and retrieves it by using `mail`, `mailx`, `mailtool`, or a similar program.

sendmail Program

The following list describes some of the capabilities of the `sendmail` program.

- `sendmail` can use different types of communications protocols, such as TCP/IP and UUCP.
- `sendmail` implements an SMTP server, message queueing, and mailing lists.
- `sendmail` controls name interpretation by using a pattern-matching system that can work with the following naming conventions.

- Domain-based naming convention. The domain technique separates the issue of physical versus logical naming. For more information on domains, refer to “Mail Addresses” on page 357.
- Improvised techniques, such as providing network names that appear local to hosts on other networks.
- Arbitrary (older) naming syntaxes.
- Disparate naming schemes.

The Solaris operating environment uses the `sendmail` program as a mail router. The following list describes some of its functions.

- `sendmail` is responsible for receiving and delivering email messages.
- `sendmail` is an interface between mail-reading programs such as `mail`, `mailx`, and `mailtool`, and mail-transport programs such as `uucp`.
- `sendmail` controls email messages that users send.
 - By evaluating the recipients’ addresses
 - By choosing an appropriate delivery program
 - By rewriting the addresses in a format that the delivery agent can handle
 - By reformatting the mail headers as required
 - By finally passing the transformed message to the mail program for delivery

For more information about the `sendmail` program, refer to the following topics.

- “`sendmail` and Its Rerouting Mechanisms” on page 372
- “`sendmail` Features” on page 374
- “`sendmail` Configuration File” on page 375

`sendmail` and Its Rerouting Mechanisms

The `sendmail` program supports three mechanisms for mail rerouting. The mechanism you choose depends on the type of change that is involved.

- A server change
- A domain-wide change
- A change for one user

Additionally, the rerouting mechanism you choose can affect the level of administration that is required. Consider the following options.

1. One rerouting mechanism is *aliasing*.

Aliasing can map names to addresses on a server-wide basis or a name service-wide basis, depending on the type of file that you use.

Consider the following advantages and disadvantages for name service aliasing.

- The use of a name service (such as NIS or NIS+) alias file permits mail rerouting changes to be administered from a single source. However, name service aliasing can create lag time when the rerouting change is propagated.

- Name service administration is usually restricted to a select group of system administrators. A normal user would not administer this file.

Consider the following advantages and disadvantages for using a server alias file.

- By using a server alias file, rerouting can be managed by anyone who can become `root` on the designated server.
- Server aliasing should create little or no lag time when the rerouting change is propagated.
- The change only affects the local server, which might be acceptable if most of the mail is sent to one server. However, if you need to propagate this change to many mail servers, use a name service.
- A normal user would not administer this change.

For more information, refer to “Mail Alias Files” on page 376 in this chapter. For a task map, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

2. The next mechanism is *forwarding*.

This mechanism permits users to administer mail rerouting. Local users can reroute their incoming mail to the following.

- Another mailbox
- A different mailer
- Another mail host

This mechanism is supported through the use of `.forward` files. For more information about these files, refer to “`.forward` Files” on page 379 in this chapter. For a task map, refer to “Administering `.forward` Files (Task Map)” on page 341 in Chapter 25.

3. The last rerouting mechanism is *inclusion*.

This mechanism allows users to maintain alias lists, instead of requiring `root` access. To provide this feature, the `root` user must create an appropriate entry in the alias file on the server. After this entry is created, the user can reroute mail as necessary. For more information on inclusion, refer to “`/etc/mail/aliases` File” on page 377 in this chapter. For a task map, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

Figure 26–3 shows how `sendmail` uses aliases. Programs that read mail, such as `/usr/bin/mailx`, can have aliases of their own, which are expanded before the message reaches `sendmail`. The aliases for `sendmail` can originate from a number of name service sources (local files, NIS or NIS+). The order of the lookup is determined by the `nsswitch.conf` file. Refer to the `nsswitch.conf(4)` man page.




FIGURE 26-3 How sendmail Uses Aliases

sendmail Features

The sendmail program provides the following features.

- sendmail is reliable. It is designed to correctly deliver every message. No message should ever become completely lost.
- sendmail uses existing software for delivery whenever possible.
- sendmail can be configured to handle complex environments, including multiple connections to a single network type, such as UUCP or Ethernet. sendmail checks the contents of an address as well as its syntax to determine which mailer to use.
- sendmail uses configuration files to control mail configuration instead of requiring that configuration information be compiled into the code.
- Users can maintain their own mailing lists. In addition, individuals can specify their own forwarding mechanism without modifying the domain-wide alias file, typically located in the domain-wide aliases that are maintained by NIS or NIS+.
- Each user can specify a custom mailer to process incoming mail, which can provide functions such as returning a message that reads: "I am on vacation." See the vacation(1) man page for more information.

- `sendmail` batches addresses to a single host to reduce network traffic.

Figure 26–4 shows how `sendmail` interacts with the other programs in the mail system.




FIGURE 26–4 Interaction of `sendmail` With Other Mail Programs

As shown in Figure 26–4, the user interacts with a mail-generating and mail-sending program. When the mail is submitted, the mail-generating program calls `sendmail`, which routes the message to the correct mailers. Because some of the senders might be network servers and some of the mailers might be network clients, you can use `sendmail` as an Internet mail gateway. See “Interactions of Mail Programs” on page 370 for a more detailed description of the process.

`sendmail` Configuration File

A *configuration file* controls the way that `sendmail` performs its functions. The configuration file determines the choice of delivery agents, address rewriting rules, and the format of the mail header.

The `sendmail` program uses the information from the `/etc/mail/sendmail.cf` file to perform its functions. Each system has a default `sendmail.cf` file that is installed in the `/etc/mail` directory. You do not need to edit or change the default configuration file for mail servers or mail clients. The only systems that require a customized configuration file are mail hosts and mail gateways.

The Solaris operating environment provides three default configuration files in the `/etc/mail` directory.

1. A configuration file that is named `main.cf` for the system (or systems) you designate as the mail host or a mail gateway

2. A configuration file that is named `subsidiary.cf`, which is a duplicate copy of the default `sendmail.cf` file
3. A configuration file that is named `submit.cf`, which is used to run `sendmail` in mail submission program mode, instead of daemon mode. For more information, refer to “New Configuration File, `submit.cf`” on page 390.

The configuration file you use on a system depends on the role of the system in your mail service.

- For mail clients or mail servers, you do not need to do anything to set up or edit the default configuration file.
- To set up a mail host or mail gateway, copy the `main.cf` file and rename it to `sendmail.cf` in the `/etc/mail` directory. Then reconfigure the `sendmail` configuration file to set the relay mailer and relay host parameters that are needed for your mail configuration. For task information, refer to “Setting Up Mail Services (Task Map)” on page 314 or “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

The following list describes some configuration parameters you can change, depending on the requirements of your site.

- Time values, which specify the following information.
 - Read timeouts. Refer to “Changes to the Timeout Option” on page 404.
 - Length of time a message remains undelivered in the queue before it is returned to the sender. Refer to “New Queue Features” on page 414. For a task map, refer to “Administering the Queue Directories (Task Map)” on page 338.
- Delivery modes, which specify how quickly mail is delivered.
- Load limits, which increase efficiency during busy periods by not attempting to deliver large messages, messages to many recipients, and messages to sites that have been down for a long time.
- Log level, which specifies the kinds of problems that are logged.

Mail Alias Files

You can use any of the following files, maps, or tables to maintain aliases.

- “.mailrc Aliases” on page 377
- “`/etc/mail/aliases` File” on page 377
- “NIS aliases Map” on page 378
- “NIS+ `mail_aliases` Table” on page 379

Your method of maintaining aliases depends on who uses the alias and who needs to be able to change the alias. Each type of alias has unique format requirements.

If you are looking for task information, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

.mailrc Aliases

Aliases that are listed in a `.mailrc` file are accessible only by the user who owns the file. This restriction enables users to establish an alias file they control and that is usable only by its owner. Aliases in a `.mailrc` file adhere to the following format.

```
alias aliasname value value value ...
```

aliasname is the name that the user uses when sending mail, and *value* is a valid email address.

If a user establishes a personal alias for `scott` that does not match the email address for `scott` in the name service, mail is routed to the wrong person when people try to reply to mail that is generated by this user. The only workaround is to use any of the other aliasing mechanisms.

/etc/mail/aliases File

Any alias that is established in the `/etc/mail/aliases` file can be used by any user who knows the name of the alias and the host name of the system that contains the file. Distribution list formats in a local `/etc/mail/aliases` file adhere to the following format.

```
aliasname: value,value,value ...
```

aliasname is the name that the user uses when sending mail to this alias, and *value* is a valid email address.

If your network is not running a name service, the `/etc/mail/aliases` file of each system should contain entries for all mail clients. You can either edit the file on each system or edit the file on one system and copy it to each of the other systems.

The aliases in the `/etc/mail/aliases` file are stored in text form. When you edit the `/etc/mail/aliases` file, you need to run the `newaliases` program to recompile the database and make the aliases available in binary form to the `sendmail` program. For task information, refer to “How to Set Up a Local Mail Alias File” on page 333 in Chapter 25. Otherwise, you can use the AdminTool’s Database Manager to administer the mail aliases that are stored in the local `/etc` files.

You can create aliases for only local names—a current host name or no host name. For example, an alias entry for user `ignatz` who has a mailbox on the system `saturn` would have the following entry in the `/etc/mail/aliases` file.

```
ignatz: ignatz@saturn
```

You should create an administrative account for each mail server. You create such an account by assigning a mailbox on the mail server to `root` and by adding an entry for `root` to the `/etc/mail/aliases` file. For example, if the system `saturn` is a mailbox server, add the entry `root: sysadmin@saturn` to the `/etc/mail/aliases` file.

Normally, only the `root` user can edit this file. However, when you use the Administration Tool, all users in group 14, which is the `sysadmin` group, can change the local file. Another option is to create the following entry.

```
aliasname: :include:/path/aliasfile
```

aliasname is the name that the user uses when sending mail, and */path/aliasfile* is the full path to the file that contains the alias list. The alias file should include email entries, one entry on each line, and no other notations.

```
user1@host1
```

```
user2@host2
```

You can define additional mail files in `/etc/mail/aliases` to keep a log or a backup copy. The following entry stores all mail that is sent to *aliasname* in *filename*.

```
aliasname: /home/backup/filename
```

You can also route the mail to another process. The following example stores a copy of the mail message in *filename* and prints a copy.

```
aliasname: "|tee -a /home/backup/filename |lp"
```

For a task map, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

NIS aliases Map

All users in a local domain can use the entries that are in the NIS `aliases` map, because the `sendmail` program can use the NIS `aliases` map instead of the local `/etc/mail/aliases` files to determine mailing addresses. For more information, refer to the `nsswitch.conf(4)` man page.

Aliases in the NIS `aliases` map adhere to the following format.

```
aliasname: value,value,value ...
```

aliasname is the name that the user uses when sending mail, and *value* is a valid email address.

The NIS `aliases` map should contain entries for all mail clients. In general, only the `root` user on the NIS master can change these entries. This type of alias might not be a good choice for aliases that are constantly changing, but such aliases can be useful if the aliases point to another alias file, as in the following syntax example.

```
aliasname: aliasname@host
```

aliasname is the name that the user uses when sending mail, and *host* is the host name for the server that contains an `/etc/mail/alias` file.

For task information, refer to “How to Set Up an NIS `mail.aliases` Map” on page 332 in Chapter 25.

NIS+ mail_aliases Table

The NIS+ mail_aliases table contains the names by which a system or person is known in the local domain. The sendmail program can use the NIS+ mail_aliases table, instead of the local /etc/mail/aliases files, to determine mailing addresses. Refer to the aliasadm(1M) and nsswitch.conf(4) man pages for more information.

Aliases in the NIS+ mail_aliases table adhere to the following format:

```
alias: expansion # ["options " # "comments"]
```

Table 26–12 describes the four columns that are in an NIS+ mail_aliases table.

TABLE 26–12 Columns in the NIS+ mail_aliases Table

Column	Description
alias	The name of the alias
expansion	The value of the alias or a list of aliases as it would appear in a sendmail /etc/mail/aliases file
options	The column that is reserved for future use
comments	The column for comments about an individual alias

The NIS+ mail_aliases table should contain entries for all mail clients. You can list, create, modify, and delete entries in the NIS+ aliases table with the aliasadm command. To use the aliasadm command, you must be a member of the NIS+ group that owns the aliases table. For task information, refer to “How to Manage Alias Entries in an NIS+ mail_aliases Table” on page 327 in Chapter 25. Alternately, you can use the AdminTool’s Database Manager to administer the NIS+ mail aliases.

Note – If you are creating a new NIS+ aliases table, you must initialize the table before you create the entries. If the table exists, no initialization is needed.

.forward Files

Users can create a .forward file in their home directories that sendmail, along with other programs, can use to redirect mail or send mail. Refer to the following topics.

- “Situations to Avoid” on page 380
- “Controls for .forward files” on page 380
- “.forward.hostname File” on page 380
- “.forward+detail File” on page 381

For a task map, refer to “Administering .forward Files (Task Map)” on page 341 in Chapter 25.

Situations to Avoid

The following list describes some situations that you can avoid or easily fix.

- If mail is not being delivered to the expected address, check the user’s .forward file. The user might have put the .forward file in the home directory of host1, which forwards mail to user@host2. When the mail arrives at host2, sendmail checks for user in the NIS or NIS+ aliases and sends the message back to user@host1, which results in a loop and more bounced mail.
- To avoid security problems, never put .forward files in the root and bin accounts. If necessary, forward the mail by using the aliases file instead.

Controls for .forward files

For the .forward files to be an effective part of mail delivery, ensure that the following controls (mostly permissions settings) are correctly applied.

- The .forward file must be writable only by the owner of the file. This restriction prevents other users from breaking security.
- The paths that lead to the home directory must be owned and writable by root only. For example, if a .forward file is in /export/home/terry, /export and /export/home must be owned and writable by root only.
- The actual home directory should be writable only by the user.
- The .forward file cannot be a symbolic link, and it cannot have more than one hard link.

.forward.hostname File

You can create a .forward.hostname file to redirect mail that is sent to a specific host. For example, if a user’s alias has changed from sandy@phoenix.example.com to sandy@example.com, place a .forward.phoenix file in the home directory for sandy.

```
% cat .forward.phoenix
sandy@example.com
"/usr/bin/vacation sandy"
% cat .vacation.msg
From: sandy@example.com (via the vacation program)
Subject: my alias has changed
```

```
My alias has changed to sandy@example.com.
Please use this alias in the future.
```

Composed September 5, 2002

The mail that I just received from you
has been forwarded to my new address.

Sandy

In this example, mail can be forwarded to the correct place while the sender is notified of the alias change. Because the `vacation` program permits only one message file, you can forward only one message at a time. However, if the message is not host specific, one vacation message file can be used by `.forward` files for many hosts.

`.forward+detail` File

Another extension to the forwarding mechanism is the `.forward+detail` file. The *detail* string can be any sequence of characters except operator characters. The operator characters are `. : % & ! ^ [ ] +`. By using this type of file, you can determine if someone else is using your email address without your knowledge. For instance, if a user tells someone to use the email address `sandy+test1@example.com`, the user would be able to identify any future mail that was delivered to this alias. By default, any mail that is sent to the `sandy+test1@example.com` alias is checked against the alias and the `.forward+detail` files. If no matches are made, the mail falls back to delivery to `sandy@example.com`, but the user is able to see a change in the `To:` mail header.

`/etc/default/sendmail` File

This file is used to store startup options for `sendmail` so that they are not removed when a host is upgraded. The following variables can be used.

`CLIENTOPTIONS="string"`

Selects additional options to be used with the client daemon, which looks in the client-only queue (`/var/spool/clientmqueue`) and acts as a client queue runner. No syntax checking is done, so be careful when making changes to this variable.

`CLIENTQUEUEINTERVAL=#`

Similar to the `QUEUEINTERVAL` option, `CLIENTQUEUEINTERVAL` sets the time interval for mail queue runs. However, the `CLIENTQUEUEINTERVAL` option controls the functions of the client daemon, instead of the master daemon.

Typically, the master daemon is able to deliver all messages to the SMTP port. However, if the message load is too high or the master daemon is not running, then messages go into the client-only queue, `/var/spool/clientmqueue`. The client daemon, which checks in the client-only queue, then acts as a client queue processor.

`ETRN_HOSTS="string"`

Enables an SMTP client and server to interact immediately without waiting for the periodic queue run intervals. The server can immediately deliver the portion of its queue that goes to the specified hosts. For more information, refer to the `etrn(1M)` man page.

MODE=`bd`

Selects the mode to start `sendmail` with. Use the `-bd` option or leave it undefined.

OPTIONS=*string*

Selects additional options to be used with the master daemon. No syntax checking is done, so be careful when making changes to this variable.

QUEUEINTERVAL=*#*

Sets the interval for mail queue runs on the master daemon. *#* can be a positive integer that is followed by either *s* for seconds, *m* for minutes, *h* for hours, *d* for days, or *w* for weeks. The syntax is checked before `sendmail` is started. If the interval is negative or if the entry does not end with an appropriate letter, the interval is ignored and `sendmail` starts with a queue interval of 15 minutes.

QUEUEOPTIONS=*p*

Enables one persistent queue runner that sleeps between queue run intervals, instead of a new queue runner for each queue run interval. You can set this option to *p*, which is the only setting available. Otherwise, this option is not set.

Mail Addresses and Mail Routing

The path a mail message follows during delivery depends on the setup of the client system and the topology of the mail domain. Each additional level of mail hosts or mail domains can add another alias resolution, but the routing process is basically the same on most hosts.

You can set up a client system to receive mail locally. Receiving mail locally is known as running `sendmail` in local mode. Local mode is the default for all mail servers and some clients. On a mail server or a mail client in local mode, a mail message is routed the following way.

Note – The following example assumes that you are using the default rule set in the `sendmail.cf` file.

1. **Expand the mail alias, if possible, and restart the local routing process.**

The mail address is expanded by checking for the mail alias in the name service and substituting the new value, if one is found. This new alias is then checked again.

2. **If the mail is local, deliver it to `/usr/lib/mail.local`.**

The mail is delivered to a local mailbox.

3. **If the mail address includes a host in this mail domain, deliver the mail to that host.**

4. If the address does not include a host in this domain, forward the mail to the mail host.

The mail host uses the same routing process as the mail server, but the mail host can receive mail that is addressed to the domain name as well as to the host name.

Interactions of sendmail With Name Services

This section describes domain names as they apply to sendmail and name services, as well as the rules for effective use of name services, and the specific interactions of sendmail with name services. For details, refer to the following topics.

- “sendmail.cf and Mail Domains” on page 383
- “sendmail and Name Services” on page 384
- “Interactions of NIS and sendmail” on page 385
- “Interactions of sendmail With NIS and DNS” on page 386
- “Interactions of NIS+ and sendmail” on page 386
- “Interactions of sendmail With NIS+ and DNS” on page 387

If you are looking for related task information, refer to “How to Use DNS With sendmail” on page 322 or “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

sendmail.cf and Mail Domains

The standard `sendmail.cf` file uses mail domains to determine whether mail is delivered directly or through a mail host. Intradomain mail is delivered through a direct SMTP connection, while interdomain mail is forwarded to a mail host.

In a secure network, only a few selected hosts are authorized to generate packets that are targeted to external destinations. Even if a host has the IP address of the remote host that is external to the mail domain, the establishment of an SMTP connection is not guaranteed. The standard `sendmail.cf` assumes the following.

- The current host is not authorized to send packets directly to a host outside the mail domain.
- The mail host is capable of forwarding the mail to an authorized host that can transmit packets directly to an external host. Actually, the mail host can itself be an authorized host.

With these assumptions, the mail host is responsible for delivering or forwarding interdomain mail.

sendmail and Name Services

sendmail imposes various requirements on name services. To improve your understanding of these requirements, this section first describes the relationship of mail domains to name service domains. Then the section describes the various requirements. Refer to the following.

- “Mail Domains and Name Service Domains” on page 384
- “Host Name Service Data” on page 384
- Man pages for `in.named(1M)`, `nis+(1)`, `nisaddent(1M)`, and `nsswitch.conf(4)`

Mail Domains and Name Service Domains

The mail domain name must be a suffix of the name service domain. For example, if the domain name of the name service is `A.B.C.D`, the mail domain name could be one of the following.

- `A.B.C.D`
- `B.C.D`
- `C.D`
- `D`

When first established, the mail domain name is often identical to the name service domain. As the network grows, the name service domain can be divided into smaller pieces to make the name service more manageable. However, the mail domain often remains undivided to provide consistent aliasing.

Host Name Service Data

This section describes the requirements that sendmail imposes on name services.

A host table or map in a name service must be set up to support three types of `gethostbyname()` queries.

- `mailhost` – Some name service configurations satisfy this requirement automatically.
- Full host name (for example, `smith.admin.acme.com`) – Many name service configurations satisfy this requirement.
- Short host name (for example, `smith`) – sendmail must connect to the mail host in order to forward external mail. To determine if a mail address is within the current mail domain, `gethostbyname()` is invoked with the full host name. If the entry is found, the address is considered internal.

NIS, NIS+, and DNS support `gethostbyname()` with a short host name as an argument, so this requirement is automatically satisfied.

Two additional rules about the host-name service need to be followed to establish efficient sendmail services within a name service.

- `gethostbyname()` with full host-name argument and short host-name argument should yield consistent results. For example, `gethostbyname(smith.admin.acme.com)` should return the same result as `gethostbyname(smith)`, as long as both functions are called from the mail domain `admin.acme.com`.
- For all name service domains under a common mail domain, `gethostbyname()` with a short host name should yield the same result. For example, if the mail domain `smith.admin.acme.com` is given, `gethostbyname(smith)` should return the same result when the call originates from either the `ebb.admin.acme.com` domain or the `esg.admin.acme.com` domain. The mail domain name is usually shorter than the name service domain, which gives this requirement special implications for various name services.

For more information about the `gethostbyname()` function, refer to the `gethostbyname(3NSL)` man page.

Interactions of NIS and `sendmail`

The following list describes the interactions of `sendmail` and NIS and provides some guidance.

- **Mail domain name** – If you are setting up NIS as the primary name service, `sendmail` automatically strips the first component of the NIS domain name and uses the result as the mail domain name. For example, `ebs.admin.acme.com` becomes `admin.acme.com`.
- **Mail host name** – You must have a `mailhost` entry in the NIS host map.
- **Full host names** – The normal NIS setup does not “understand” the full host name. Rather than trying to make NIS understand the full host name, turn off this requirement from the `sendmail` side by editing the `sendmail.cf` file and replacing all occurrences of `%l` with `%y`. This change turns off `sendmail`’s interdomain mail detection. If the target host can be resolved to an IP address, a direct SMTP delivery is attempted. Ensure that your NIS host map does not contain any host entry that is external to the current mail domain. Otherwise, you need to further customize the `sendmail.cf` file.
- **Matching full and short host names** – Follow the previous instructions on how to turn off `gethostbyname()` for a full host name.
- **Multiple NIS domains in one mail domain** – All NIS host maps under a common mail domain should have the same set of host entries. For example, the host map in the `ebs.admin.acme.com` domain should be the same as the host map in the `esg.admin.acme.com`. Otherwise, one address might work in one NIS domain, but fail in the other NIS domain.

For task information, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

Interactions of sendmail With NIS and DNS

The following list describes the interactions of sendmail with NIS and DNS and provides some guidance.

- **Mail domain name** – If you are setting up NIS as the primary name service, sendmail automatically strips the first component of the NIS domain name and uses the result as the mail domain name. For example, `ebs.admin.acme.com` becomes `admin.acme.com`.
- **Mail host name** – When the DNS forwarding feature is turned on, queries that NIS cannot resolve are forwarded to DNS, so you do not need a `mailhost` entry in the NIS host map.
- **Full host names** – Although NIS does not “understand” full host names, DNS does understand. This requirement is satisfied when you follow the regular procedure for setting up NIS and DNS.
- **Matching full and short host names** – For every host entry in the NIS host table, you must have a corresponding host entry in DNS.
- **Multiple NIS domains in one mail domain** – All NIS host maps under a common mail domain should have the same set of host entries. For example, the host map in the `ebs.admin.acme.com` domain should be the same as the host map in the `esg.admin.acme.com` domain. Otherwise, one address might work in one NIS domain, but fail in the other NIS domain.

For task information, refer to “How to Use DNS With sendmail” on page 322 and “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

Interactions of NIS+ and sendmail

The following list describes the interactions of sendmail with NIS+ and provides some guidance.

- **Mail domain name** – If you are setting up NIS+ as your primary name service, sendmail can check the mail domain from the NIS+ `sendmailvars` table, a two-column NIS+ table with one key column and one value column. To set up your mail domain, you must add one entry to this table. This entry should have the key column set to the literal string `maildomain` and the value column set to your mail domain name (for example, `admin.acme.com`). Although NIS+ allows any string in the `sendmailvars` table, the suffix rule still applies for the mail system to work correctly. You can use `nistbladm` to add the `maildomain` entry to the `sendmailvars` table. Notice in the following example that the mail domain is a suffix of the NIS+ domain.

```
nistbladm -A key="maildomain" value=<mail domain> sendmailvars.org_dir.<NIS+ domain>
```

- **Mailhost host name** – You must have a `mailhost` entry in the NIS+ hosts table.
- **Full host names** – NIS+ “understands” the full host name. Following the regular NIS+ setup procedure satisfies this requirement.

- **Matching full and short host names** – To satisfy this requirement, you can duplicate the entries in the host table, or you can enter all host entries in the user name service domains into a master host table at mail domain level.
- **Multiple NIS domains in one mail domain** – To satisfy this requirement, you can duplicate the entries in all the host tables, or you can type all host entries in the user name service domains into a master host table at mail domain level. In effect, you are merging (logical or physical) multiple host tables into one host table. Therefore, the same host name cannot be reused in the multiple name service domain that shares a common mail domain.

For task information, refer to “Administering Mail Alias Files (Task Map)” on page 326 in Chapter 25.

Interactions of `sendmail` With NIS+ and DNS

The following list describes the interactions of `sendmail` with NIS+ and DNS and provides some guidance.

- **Mail domain name** — If you are setting up NIS+ as your primary name service, `sendmail` can check the mail domain from the NIS+ `sendmailvars` table, a two-column NIS+ table with one key column and one value column. To set up your mail domain, you must add one entry to this table. This entry should have the key column set to the literal string `maildomain` and the value column set to the your mail domain name (for example, `admin.acme.com`). Although NIS+ allows any string in the `sendmailvars` table, the suffix rule still applies for the mail system to work correctly. You can use `nistbladm` to add the `maildomain` entry to the `sendmailvars` table. Notice in the following example that the mail domain is a suffix of the NIS+ domain.

```
nistbladm -A key="maildomain" value=<mail domain> sendmailvars.org_dir.<NIS+ domain>
```

- **Mailhost host name** — If your network uses both NIS+ and DNS as the source for the host database, you can put the `mailhost` entry in either the NIS+ or DNS host table. Ensure that your users include NIS+ and DNS as the source for the host database in the `/etc/nsswitch.conf` file.
- **Full host names** — Both NIS+ and DNS “understand” full host names. Following the regular NIS+ and DNS setup procedures satisfies this requirement.
- **Matching full and short host names** — For every host entry in the NIS+ host table, you must have a corresponding host entry in DNS.
- **Multiple NIS domains in one mail domain** — To satisfy this requirement, you can duplicate the entries in all the host tables, or you can type all host entries in the user name service domains into a master host table at the mail domain level.

For task information, refer to “Administering Mail Alias Files (Task Map)” on page 326 and “How to Use DNS With `sendmail`” on page 322 in Chapter 25.

Composed September 5, 2002

What's New With Mail Services (Reference)

Chapter 24 provides an introduction to the components of mail services and a description of a typical mail configuration. Chapter 25 explains how to set up and administer an electronic mail system with standard configuration files. Chapter 26 describes in greater detail the components of mail services. Chapter 26 also describes the mail service programs and files, the mail routing process, and the interactions of `sendmail` with name services. This chapter describes the new features that are included in version 8.12 of `sendmail`, the version that is in this Solaris 9 release. You can also read about changes to `mail.local`, `mailstats`, and `makemap`. This chapter also describes a new maintenance utility, `editmap`. The following list can help you in your search on a specific topic.

- “Changes to `sendmail`” on page 389
- “Changes to `mail.local`” on page 418
- “Changes to `mailstats`” on page 419
- “Changes to `makemap`” on page 419
- “New Command, `editmap`” on page 420
- “Other Changes and Features of Interest” on page 421

For details that are not covered in this chapter, see the man pages for `sendmail(1M)`, `mail.local(1M)`, `mailstats(1)`, `makemap(1M)`, and `editmap(1M)`.

Changes to `sendmail`

This section contains information on the following topics.

- “New Configuration File, `submit.cf`” on page 390
- “New or Deprecated Command-Line Options” on page 391
- “New and Revised Configuration File Options and Related Topics” on page 392
- “New Defined Macros for `sendmail`” on page 406
- “New and Revised m4 Configuration Macros for `sendmail`” on page 408

- “New Delivery Agent Flags” on page 412
- “New Equates for Delivery Agents” on page 413
- “New Queue Features” on page 414
- “Changes for LDAP in `sendmail`” on page 415
- “New Built-in Mailer Feature” on page 416
- “New Rule Sets” on page 416
- “Changes to Files” on page 417
- “IPv6 Addresses in Configuration” on page 418

New Configuration File, `submit.cf`

Version 8.12 of `sendmail` includes an additional configuration file, `/etc/mail/submit.cf`. This new file, `submit.cf`, is used to run `sendmail` in mail submission program mode instead of daemon mode. Mail submission program mode, unlike daemon mode, does not require `root` privilege, so this new paradigm provides better security.

See the following list of functions for `submit.cf`:

- `sendmail` uses `submit.cf` to run in mail submission program (MSP) mode, which submits email messages and can be invoked by programs (such as `mailx`), as well as by users. Refer to the description of the `-Ac` option and the `-Am` option in “New or Deprecated Command-Line Options” on page 391.
- `submit.cf` is used in the following operating modes:
 - bm, which is the default operating mode
 - bs, which uses standard input to run SMTP
 - bt, which is the test mode that is used to resolve addresses
- `sendmail`, when using `submit.cf`, does not run as an SMTP daemon.
- `sendmail`, when using `submit.cf`, uses `/var/spool/clientmqueue`, the client-only mail queue, which holds messages that were not delivered to the `sendmail` daemon. Messages in the client-only queue are delivered by the client “daemon,” which is really acting as a client queue runner.
- By default, `sendmail` uses `submit.cf` periodically to run the MSP queue (otherwise, known as the client-only queue), `/var/spool/clientmqueue`.

```
/usr/lib/sendmail -Ac -q15m
```

Note the following:

- `submit.cf` is provided automatically when you install or upgrade to the Solaris 9 operating environment.
- `submit.cf` requires no planning or preliminary procedures prior to the installation of the Solaris 9 operating environment.
- Unless you specify a configuration file, `sendmail` automatically uses `submit.cf` as required. Basically, `sendmail` knows which tasks are appropriate for `submit.cf` and which tasks are appropriate for `sendmail.cf`.

- `submit.cf` is not to be modified.

Functions That Distinguish `sendmail.cf` From `submit.cf`

The `sendmail.cf` configuration file is for the daemon mode. When using this file, `sendmail` is acting as a mail transfer agent (MTA), which is started by `root`.

```
/usr/lib/sendmail -L sm-mta -bd -qlh
```

See the following list of other distinguishing functions for `sendmail.cf`:

- By default, `sendmail.cf` accepts SMTP connections on ports 25 and 587.
- By default, `sendmail.cf` runs the main queue, `/var/spool/queue`.

Functional Changes in `sendmail`

With the addition of `submit.cf`, the following functional changes have occurred:

- In version 8.12 of `sendmail`, only `root` can run the mail queue. For further details, refer to the changes that are described in the `mailq(1)` man page. For new task information, refer to “Administering the Queue Directories (Task Map)” on page 338.
- Because the mail submission program mode runs without `root` privilege, which might prevent `sendmail` from having access to certain files (such as, the `.forward` files), the `-bv` option for `sendmail` could give the user misleading output. No workaround is available.
- Prior to `sendmail` version 8.12, if you were not running the `sendmail` daemon (that is, running in daemon mode), you would only prevent the delivery of inbound mail. Now, in `sendmail` version 8.12, if you are not running the `sendmail` daemon with the default configuration, you also prevent the delivery of outbound mail. The client queue runner (also known as the mail submission program) must be able to submit mail to the daemon on the local SMTP port. If the client queue runner tries to open an SMTP session with the local host and the daemon is not listening on the SMTP port, the mail remains in the queue. The default configuration does run a daemon, so this problem does not occur if you are using the default configuration. However, if you have disabled your daemon, refer to “Managing Mail Delivery by Using an Alternate Configuration (Task)” on page 325 for a way to resolve this problem.

New or Deprecated Command-Line Options

The following table describes new command-line options for `sendmail`. Other command-line options are described in the `sendmail(1M)` man page.

TABLE 27-1 New Command-Line Options for `sendmail`

Option	Description
-Ac	Indicates that you want to use the configuration file, <code>submit.cf</code> , even if the operation mode does not indicate an initial mail submission. For more information about <code>submit.cf</code> , refer to “New Configuration File, <code>submit.cf</code> ” on page 390.
-Am	Indicates that you want to use the configuration file, <code>sendmail.cf</code> , even if the operation mode indicates an initial mail submission. For more information, refer to “New Configuration File, <code>submit.cf</code> ” on page 390.
-bP	Indicates that you are printing the number of entries in each queue.
-G	Indicates that the message that is being submitted from the command line is for relaying, not for initial submission. The message is rejected if the addresses are not fully qualified. No canonicalization is done. As is noted in the Release Notes that are part of the <code>sendmail</code> distribution on ftp://ftp.sendmail.org , improperly formed messages might be rejected in future releases.
-L <i>tag</i>	Sets the identifier that is used for syslog messages to the supplied <i>tag</i> .
-q[!]I <i>substring</i>	Processes only jobs that contain this <i>substring</i> of one of the recipients. When <code>!</code> is added, the option processes only jobs that do not have this <i>substring</i> of one of the recipients.
-q[!]R <i>substring</i>	Processes only jobs that contain this <i>substring</i> of the queue ID. When <code>!</code> is added, the option processes only jobs that do not have this <i>substring</i> of the queue ID.
-q[!]S <i>substring</i>	Processes only jobs that contain this <i>substring</i> of the sender. When <code>!</code> is added, the option processes only jobs that do not have this <i>substring</i> of the sender.
-qf	Processes saved messages in the queue once, without using the <code>fork</code> system call, and runs the process in the foreground. Refer to the <code>fork(2)</code> man page.
-qGname	Processes only the messages in the <i>name</i> queue group.
-qp <code>time</code>	Processes saved messages in the queue at a specific interval of time with a single child that is forked for each queue. The child sleeps between queue runs. This new option is similar to the <code>-qtime</code> , which periodically forks a child to process the queue.
-U	As is noted in the Release Notes that are part of the <code>sendmail</code> distribution on ftp://ftp.sendmail.org , this option is not available in version 8.12. Mail user agents should use the <code>-G</code> argument.

New and Revised Configuration File Options and Related Topics

This section contains a table of new and revised configuration file options and information on the following related topics.

- “Deprecated and Unsupported Configuration File Options for `sendmail`” on page 400
- “New `ClientPortOptions` Option” on page 400

- “Changes to DaemonPortOptions Option” on page 401
- “Additional Arguments for the PidFile and ProcessTitlePrefix Options” on page 403
- “Changes to the PrivacyOptions Option” on page 403
- “Changes to the Timeout Option” on page 404

When you declare these options, use one of the following syntaxes.

```
O OptionName=argument      # for the configuration file
-O OptionName=argument      # for the command line
define('m4Name', argument) # for m4 configuration
```

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

The following table describes new and revised options for `sendmail`.

TABLE 27–2 New and Revised Options for `sendmail`

Option	Description
BadRcptThrottle	m4 name: <code>confBAD_RCPT_THROTTLE</code> Argument: <i>number</i> The new option limits the rate that recipients in the SMTP envelope are accepted after a threshold number of recipients has been rejected.
ClientPortOption	For details, see “New ClientPortOptions Option” on page 400.
ConnectionRateThrottle	m4 name: <code>confCONNECTION_RATE_THROTTLE</code> Argument: <i>number</i> The option <code>ConnectionRateThrottle</code> now limits the number of connections per second to each daemon, not the total number of connections.
ControlSocketName	m4 name: <code>confCONTROL_SOCKET_NAME</code> Argument: <i>filename</i>. The recommended socket name is <code>/var/spool/mqueue/.smcontrol</code>. For security, this UNIX domain socket must be in a directory that is accessible only by <code>root</code>. When it is set, this new option creates a control socket for daemon management. This option enables an external program to control and query the status of the running <code>sendmail</code> daemon by way of a named socket. The socket is similar to the <code>ctlinnd</code> interface to the INN news server. If this option is not set, no control socket is available.
DaemonPortOptions	For details, see “Changes to DaemonPortOptions Option” on page 401.

TABLE 27-2 New and Revised Options for sendmail (Continued)

Option	Description
DataFileBufferSize	m4 name: confDF_BUFFER_SIZE Argument: <i>number</i> The new option controls the maximum size (in bytes) of a memory-buffered data (df) file before a disk-based file is used. The default is 4096 bytes. You should not have to change the default for the Solaris operating environment.
DeadLetterDrop	m4 name: confDEAD_LETTER_DROP Argument: <i>filename</i> The new option, which you should not need to set, defines the location of the system-wide dead.letter file, which was formerly hard-coded to /usr/tmp/dead.letter.
DelayLA	m4 name: confDELAY_LA Argument: <i>number</i> If this new option is set to a value greater than zero, the option does the following: Delays connections by one second when the load averages exceed a specified value Delays the execution of most SMTP commands by one second Otherwise, if the option is not set, the default value, which is zero, does not change the behavior of sendmail.
DeliverByMin	m4 name: confDELIVER_BY_MIN Argument: <i>time</i> The new option enables a client to specify a minimum amount of time for an email message to be delivered, as specified in RFC 2852, Deliver By SMTP Service Extension. If the time is set to zero, no time is listed. If the time is set to less than zero, the extension is not offered. If the time is set to greater than zero, the extension is listed as the minimum time for the EHLO keyword, DELIVERBY.
DirectSubmissionModifiers	m4 name: confDIRECT_SUBMISSION_MODIFIERS Argument: <i>modifiers</i> The new option defines \${daemon_flags} for direct (command-line) submissions. If this option is not set, the value of \${daemon_flags} is either CC f, if the option -G is used, or c u.

TABLE 27-2 New and Revised Options for sendmail (Continued)

Option	Description
DontBlameSendmail	You can use the following new arguments. The argument, <code>NonRootSafeAddr</code>, has been added. When <code>sendmail</code> does not have enough privileges to run a <code>.forward</code> program or deliver to a file as the owner of that file, addresses are marked unsafe. Furthermore, if <code>RunAsUser</code> is set, users cannot use programs or deliver to files in their <code>.forward</code> programs. Use <code>NonRootSafeAddr</code> to resolve these problems.
DoubleBounceAddress	m4 name: <code>confDOUBLE_BOUNCE_ADDRESS</code> Argument: <i>address</i>. The default is <code>postmaster</code>. If an error occurs when <code>sendmail</code> is sending an error message, <code>sendmail</code> sends the "double-bounced" error message to the address that is specified by the argument to this option.
FallBackMXhost	m4 name: <code>confFALLBACK_MX</code> Argument: fully qualified domain name. This option now includes MX record lookups. To use the old behavior of no MX record lookups, you must put the name in square brackets.
FastSplit	m4 name: <code>confFAST_SPLIT</code> Argument: <i>number</i>. The default value is one. This new option does the following: If the option is set to a value greater than zero, the initial MX lookups on addresses are suppressed when they are sorted, which might result in faster envelope splitting. If the mail is submitted from the command line, the value can limit the number of processes that are used to deliver the envelopes. If more envelopes are created, they are put in the queue and must be resolved with a queue run.
LDAPDefaultSpec	m4 name: <code>confLDAP_DEFAULT_SPEC</code> Argument: Class switch with appropriate definition (for example, <code>-hhost</code>, <code>-pport</code>, <code>-dbind DN</code>). The new option allows a default map specification for LDAP maps. The assigned default settings are used for all LDAP maps unless other individual map specifications are made with the <code>K</code> command. Set this option before defining any LDAP maps.

TABLE 27-2 New and Revised Options for sendmail (Continued)

Option	Description
MailboxDatabase	m4 name: confMAILBOX_DATABASE Argument: pw, which uses getpwnam(), is the default value. No other values are supported. The new option specifies the type of mailbox database that is used to check for local recipients.
MaxHeadersLength	m4 name: confMAX_HEADERS_LENGTH Argument: <i>number</i> This option specifies a maximum length for the sum of all headers and can be used to prevent a denial-of-service attack. The default is 32768. A warning is issued if a value less than 16384 is used. You should not have to change the default value for the Solaris operating environment.
MaxMimeHeaderLength	m4 name: confMAX_MIME_HEADER_LENGTH Argument: <i>number</i> This option sets the maximum length of certain MIME header field values to <i>x</i> number of characters. Also, for parameters within headers, you can specify a maximum length of <i>y</i>. The combined values look like <i>x/y</i>. If <i>y</i> is not specified, half of <i>x</i> is used. If no values are set, the default is 0, which means no checks are made. This option is intended to protect mail user agents from buffer-overflow attacks. The suggested values are in the range of 256/128 to 1024/256. A warning is issued if values less than 128/40 are used.
MaxQueueChildren	m4 name: confMAX_QUEUE_CHILDREN Argument: <i>number</i> This new option limits the number of concurrently active queue-runner processes to the number that is specified in the argument. The option helps to limit the system resources that are used when the queue is processed. When the total number of queue runners for multiple queue groups exceeds the defined argument, the remaining queue groups are run later.
MaxRecipientsPerMessage	m4 name: confMAX_RCPTS_PER_MESSAGE Argument: <i>number</i> If it is set, this option allows no more than the specified number of recipients in an SMTP envelope. The minimum argument is 100. You can still declare this option from both the command line and the configuration file. However, normal users can now set it from the command line to enable the override of messages that are submitted through sendmail -bs. In this instance, sendmail does not relinquish its root privileges.

TABLE 27-2 New and Revised Options for sendmail (Continued)

Option	Description
MaxRunnersPerQueue	m4 name: confMAX_RUNNERS_PER_QUEUE Argument: <i>number</i>. The default is one. Consider your resources carefully and do not set this value too high. This new option specifies the maximum number of queue runners per queue group. The queue runners work in parallel on a queue group's messages, which is useful when the processing of a message might delay the processing of subsequent messages.
NiceQueueRun	m4 name: confNICE_QUEUE_RUN Argument: <i>number</i> This new option sets the priority of queue runners. Refer to the <code>nice(1)</code> man page.
PidFile	m4 name: confPID_file Argument: See "Additional Arguments for the PidFile and ProcessTitlePrefix Options" on page 403. This new option defines the location of the pid file. The file name is macro expanded before it is opened. The default is <code>/var/run/sendmail.pid</code>.
PrivacyOptions	For details, see "Changes to the PrivacyOptions Option" on page 403.
ProcessTitlePrefix	m4 name: confPROCESS_TITLE_PREFIX Argument: See "Additional Arguments for the PidFile and ProcessTitlePrefix Options" on page 403. The new option specifies a prefix string for the process title that is shown in <code>/usr/ucb/ps auxww</code> listings. The string is macro processed. You should not have to make any changes for the Solaris operating environment.
QueueFileMode	m4 name: confQUEUE_FILE_MODE Argument: <i>number</i> This new option provides the default permissions in octal for queue files. If this option is not set, <code>sendmail</code> uses <code>0600</code>. However, if the option's real and effective user ID is different, <code>sendmail</code> uses <code>0644</code>.
QueueLA	m4 name: confQUEUE_LA Argument: <i>number</i> The default value has changed from eight to eight times the number of processors online when the system starts. For single-processor machines, this change has no effect. Changing this value overrides the default and prevents the number of processors from being considered. Therefore, the effect of any value changes should be well understood.

TABLE 27-2 New and Revised Options for sendmail (Continued)

Option	Description
QueueSortOrder	m4 name: confQUEUE_SORT_ORDER This option sets the algorithm that is used for sorting the queue. The default value is <i>priority</i>, which sorts the queue by message priority. Note the following changes. The <i>host</i> argument now reverses the host name before sorting, which means domains are grouped to run through the queue together. This improvement provides better opportunities for use of the connection cache, if available. The new <i>filename</i> argument sorts the queue by file name, which avoids the opening and reading of each queue file when preparing to run the queue. The new <i>modification</i> argument sorts the queue by time of modification, starting with the oldest entries of the <i>qf</i> file. The new <i>random</i> argument sorts the queue randomly, which avoids contention, if several queue runners have manually been started. For more information, refer to QueueSortOrder in the sendmail(1M) man page.
RefuseLA	m4 name: confREFUSE_LA Argument: <i>number</i> The default value has changed from 12 to 12 times the number of processors online when the system starts. For single-processor machines, this change has no effect. A change of this value overrides the default and prevents the number of processors from being considered. Therefore, the effect of any value changes should be well understood.
ResolverOptions	Two changes have been made. When attempting to canonify a host name, some name servers that are down return a temporary failure message, <i>SERVFAIL</i>, for IPv6 <i>T_AAAA</i> lookups. You can use this new argument, <i>WorkAroundBrokenAAAA</i>, to avoid this behavior. Also, the <i>RES_USE_INET6</i> argument is controlled by a new flag, <i>use_inet6</i>. For more information, refer to the <i>resolver(3RESOLV)</i> man page.
RrtImpliesDsn	m4 name: confRRT_IMPLIES_DSN Argument: <i>true</i> or <i>false</i> If the new option is set, a "Return-Receipt-To:" header causes the request of a delivery status notification (DSN), which is sent to the envelope sender, not to the address that is specified in the header.

TABLE 27–2 New and Revised Options for sendmail (Continued)

Option	Description
SendMimeErrors	m4 name: <code>confMIME_FORMAT_ERRORS</code> Argument: <code>true</code> or <code>false</code> The default is now <code>true</code>.
SharedMemoryKey	m4 name: <code>confSHARED_MEMORY_KEY</code> Argument: <i>number</i> This new option permits you to use shared memory, if it is available, to store free space for queue file systems. This option minimizes the number of system calls to check for available space.
SuperSafe	m4 name: <code>confSAFE_QUEUE</code> Argument: <code>true</code>, <code>false</code>, or <code>interactive</code>. The default and recommended value is <code>true</code>. Avoid using <code>false</code>. If this option is set to <code>true</code>, the queue file is always instantiated, even if you are attempting immediate delivery. You can use the <code>interactive</code> value together with <code>DeliveryMode=i</code> to skip some synchronization calls that are doubled in the code execution path for this mode.
Timeout	For details, see “Changes to the Timeout Option” on page 404.
TrustedUser	m4 name: <code>confTRUSTED_USER</code> Argument: <i>user name</i> or <i>user numeric ID</i> The new option enables you to specify a user name (instead of <code>root</code>) to own important files. If this option is set, generated alias databases and the control socket—if configured—are automatically owned by this user. This option requires <code>HASFCHOWN</code>. For information about <code>HASFCHOWN</code>, see “Flags Used and Not Used to Compile sendmail” on page 352. Only <code>TrustedUser</code>, <code>root</code>, and class <code>t</code> (<code>\$=t</code>) users can rebuild the alias map.
UseMSP	m4 name: <code>confUSE_MSP</code> Argument: <code>true</code> or <code>false</code>. The default is <code>false</code>. This new option permits group-writable queue files, if the group is the same as that of a set-group-id sendmail binary. In <code>submit.cf</code>, this option must be set to <code>true</code>.
XscriptFileBufferSize	m4 name: <code>confXF_BUFFER_SIZE</code> Argument: <i>number</i> The new option controls the maximum size (in bytes) of a memory-buffered transcript (<code>xf</code>) file before a disk-based file is used. The default is 4096 bytes. You should not have to change this default for the Solaris operating environment.

Deprecated and Unsupported Configuration File Options for sendmail

Refer to the following table for a list of deprecated configuration file options. The table includes the `AutoRebuildAliases` option, which is not in version 8.12 of `sendmail`.

TABLE 27-3 Deprecated and Unsupported Configuration File Options for `sendmail`

Option	Description
<code>AutoRebuildAliases</code>	Because a denial-of-service attack could occur if this option is set, this option is not in version 8.12 of <code>sendmail</code> . Refer to the Release Notes that are part of the <code>sendmail</code> distribution at ftp://ftp.sendmail.org . A user could kill the <code>sendmail</code> process while the aliases file is being rebuilt and leave the file in an inconsistent state. Furthermore, because <code>AutoRebuildAliases</code> is not available, <code>newaliases</code> must be run manually now in order for changes to <code>/etc/mail/aliases</code> to become effective. Also, you must remember that because <code>sendmail</code> is no longer <code>setuid root</code> , only <code>root</code> can run <code>newaliases</code> .
<code>MeToo</code>	This option, which now defaults to <code>True</code> , has been deprecated. Refer to the Release Notes that are part of the <code>sendmail</code> distribution at ftp://ftp.sendmail.org .
<code>UnsafeGroupWrites</code>	This option is deprecated. If required, you should now use the <code>GroupWritableForwardFileSafe</code> and <code>GroupWritableIncludeFileSafe</code> arguments for the <code>DontBlameSendmail</code> option.
<code>UseErrorsTo</code>	This option is deprecated. Furthermore, because this option violates RFC 1123, you should avoid using it.

New ClientPortOptions Option

The new `ClientPortOptions` option is for outgoing connections and is similar to the `DaemonPortOptions` option. This option sets the client SMTP options, which are a sequence of *key=value* pairs. To declare this option, use one of the following syntaxes. For formatting purposes, the example includes two pairs. However, you can apply one or more pairs.

```
O ClientPortOptions=pair, pair          # for the configuration file
-O ClientPortOptions=pair, pair          # for the command line
define(`confCLIENT_OPTIONS', `pair, pair') # for m4 configuration
```

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

The following table describes the new keys for this option.

TABLE 27–4 New Keys for ClientPortOptions

Key	Description
Addr	Specifies the address mask. The value can be a numeric address in dot notation or a network name. If the pair is omitted, the default is INADDR_ANY, which accepts connections from any network.
Family	Specifies the address family. The key's default is <code>inet</code> for AF_INET. Other values are <code>inet6</code> for AF_INET6, <code>iso</code> for AF_ISO, <code>ns</code> for AF_NS, and <code>x.25</code> for AF_CCITT.
Listen	Specifies the size of the listen queue. The key defaults to 10. You should not have to change this default for the Solaris operating environment.
Port	Specifies the name and number of the listening port. The key defaults to <code>smtp</code> .
RcvBufSize	Specifies the size of the TCP/IP send buffer. The key has no default value, which means that no size specifications are automatically made. If the option is set to a value greater than zero, that value is used. You should not have to limit the size of this buffer for the Solaris operating environment.
Modifier	Specifies flags for <code>sendmail</code>: The <code>h</code> flag uses the name that corresponds to the outgoing interface address for the HELO or EHLO commands, whether it was chosen by the connection parameter or by the default. The <code>A</code> flag disables AUTH. This flag can also be used with the <code>Modifier</code> key for <code>DaemonPortOptions</code>. Refer to “Changes to <code>DaemonPortOptions</code> Option” on page 401. The <code>S</code> flag turns off the use of or the offer to use STARTTLS when email is being delivered or received.

Changes to DaemonPortOptions Option

The following tables describe the new features.

- Table 27–5
- Table 27–6

To declare this option, use one of the following syntaxes. In the example, *pair* refers to *key=value*. For formatting purposes, the example includes two pairs. However, you can apply one or more pairs.

```
O DaemonPortOptions=pair, pair           # for the configuration file
-ODaemonPortOptions=pair, pair          # for the command line
define('confDAEMON_OPTIONS', 'pair, pair') # for m4 configuration
```

Note – To avoid security risks, `sendmail` relinquishes its root permissions when you set this option from the command line.

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

The following table describes new and revised keys for the `DaemonPortOptions` option.

TABLE 27–5 New and Revised Keys for `DaemonPortOptions`

Key	Description
Name	A new key that specifies a user-definable name for <code>sendmail</code> . This key is used for error messages and for logging. The default is <code>MTA</code> .
Modifier	A new key that specifies values for <code>sendmail</code> that can be listed in a sequence without delimiters. For a list of values, see Table 27–6.
Family	Unless a <code>Family</code> is specified in a <code>DaemonPortOptions</code> option, <code>inet</code> is now the only default. If IPv6 users also want to listen on IPv6 interfaces, they can configure additional sockets into <code>sendmail.cf</code> by adding a <code>Family=inet6</code> setting to a <code>DaemonPortOptions</code> option.

The following table describes the values for the new `Modifier` key.

TABLE 27–6 Values for the New `Modifier` Key

Value	Description
A	Disables AUTH by overriding the <code>Modifier</code> value of <code>a</code> . Can be used with the <code>Modifier</code> key for <code>ClientPortOptions</code> . Refer to “New <code>ClientPortOptions</code> Option” on page 400.
C	Does not perform host-name canonification.
E	Disallows the ETRN command.
O	Ignores the socket if a failure should occur.
S	Turns off the use or the offer to use STARTTLS when email is being delivered or received. Can be used with the <code>Modifier</code> key for <code>ClientPortOptions</code> .
a	Requires authentication.
b	Binds to the interface that receives the mail.

TABLE 27-6 Values for the New Modifier Key (Continued)

Value	Description
c	Performs host-name canonification. Use this value only in configuration file declarations.
f	Requires fully qualified host names. Use this value only in configuration file declarations.
h	Uses the interface's name for the outgoing HELO command.
u	Allows unqualified addresses. Use this value only in configuration file declarations.

Additional Arguments for the PidFile and ProcessTitlePrefix Options

The following table describes additional macro-processed arguments for the PidFile and ProcessTitlePrefix options. For more information about these options, see Table 27-2.

TABLE 27-7 Arguments for the PidFile and ProcessTitlePrefix Options

Macro	Description
<code>\${daemon_addr}</code>	Provides daemon address (for example, 0.0.0.0)
<code>\${daemon_family}</code>	Provides daemon family (for example, inet, and inet6)
<code>\${daemon_info}</code>	Provides daemon information (for example, SMTP+queueing@00:30:00)
<code>\${daemon_name}</code>	Provides daemon name (for example, MSA)
<code>\${daemon_port}</code>	Provides daemon port (for example, 25)
<code>\${queue_interval}</code>	Provides queue run interval (for example, 00:30:00)

Changes to the PrivacyOptions Option

New and revised arguments for PrivacyOptions (popt) are described in the following table. You can declare this option from the command line without sendmail relinquishing its root privilege. To declare this sendmail option, use one of the following syntaxes.

```
O PrivacyOptions=argument          # for the configuration file
-O PrivacyOptions=argument          # for the command line
define('confPRIVACY_FLAGS', 'argument') # for m4 configuration
```

If you need to build a new sendmail.cf file, refer to “Building the sendmail.cf Configuration File (Task)” on page 323, in Chapter 25.

The following table provides descriptions of new and revised arguments for the PrivacyOptions option.

TABLE 27-8 New and Revised Arguments for PrivacyOptions

Argument	Description
goaway	This argument no longer accepts the following flags: noetrn, restrictmailq, restrictqrun, restrictexpand, nobodyreturn, and noreceipts.
nobodyreturn	This argument instructs sendmail not to include the body of the original message in delivery status notifications.
noreceipts	When this argument is set, delivery status notification (DSN) is not announced.
restrictexpand	This argument instructs sendmail to drop privileges when the -bv option is given by users who are neither root nor TrustedUser. The users cannot read private aliases, .forward files, or :include: files. This argument also overrides the -v command-line option.

Changes to the Timeout Option

The following table provides information about the changes to the Timeout option. Specifically, this sendmail option has some new keywords and a new value for ident. In the Solaris operating environment, you should not need to change the default values for the keywords that are listed in the table. However, if you choose to make a change, use the *keyword=value* syntax. The *value* is a time interval. Refer to the following examples.

```
O Timeout.keyword=value    # for the configuration file
-OTimeout.keyword=value    # for the command line
define('m4_name', value)  # for m4 configuration
```

If you need to build a new sendmail.cf file, refer to “Building the sendmail.cf Configuration File (Task)” on page 323 in Chapter 25.

Note – To avoid security risks, sendmail relinquishes its root permissions when you set this option from the command line.

TABLE 27–9 New and Revised Settings for Timeout

Keyword	Default Value	Description
<code>aconnect</code>	0	m4 name: <code>confTO_ACONNECT</code> Limits the total time to wait for all connections to succeed for a single delivery attempt. The maximum value is unspecified.
<code>control</code>	2m	m4 name: <code>confTO_CONTROL</code> Limits the total time that is dedicated to completing a control socket request.
<code>ident</code>	5s	m4 name: <code>confTO_IDENT</code> Defaults to 5 seconds—instead of 30 seconds—to prevent the common delays that are associated with mailing to a site that drops <code>IDENT</code> packets. No maximum value is specified.
<code>lhlo</code>	2m	m4 name: <code>confTO_LHLO</code> Limits the time to wait for a reply from an LMTP <code>LHLO</code> command. No maximum value is specified.
<code>queuereturn</code>	5d	m4 name: <code>confTO_QUEUERETURN</code> Includes the value <code>now</code> , which immediately bounces entries from the queue without a delivery attempt.
<code>resolver.retrans</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRANS</code> Sets the resolver's retransmission time interval (in seconds), which applies to <code>resolver.retrans.first</code> and <code>resolver.retrans.normal</code> .
<code>resolver.retrans.first</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRANS_FIRST</code> Sets the resolver's retransmission time interval (in seconds) for the first attempt to deliver a message.
<code>resolver.retrans.normal</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRANS_NORMAL</code> Sets the resolver's retransmission time interval (in seconds) for all resolver lookups, except the first delivery attempt.
<code>resolver.retry</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRY</code> Sets the number of times to retransmit a resolver query, which applies to <code>Timeout.resolver.retry.first</code> and <code>Timeout.resolver.retry.normal</code> .
<code>resolver.retry.first</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRY_FIRST</code> Sets the number of times to retransmit a resolver query for the first attempt to deliver a message.

TABLE 27-9 New and Revised Settings for Timeout (Continued)

Keyword	Default Value	Description
<code>resolver.retry.normal</code>	<i>varies</i>	m4 name: <code>confTO_RESOLVER_RETRY_NORMAL</code> Sets the number of times to retransmit a resolver query for all resolver lookups, except the first delivery attempt.

New Defined Macros for `sendmail`

The following table describes new macros that are reserved for use by the `sendmail` program. The macros' values are assigned internally. For more information, refer to the `sendmail(1M)` man page.

TABLE 27-10 Defined Macros for `sendmail`

Macro	Description
<code>\${addr_type}</code>	Identifies the current address as an envelope sender or a recipient address.
<code>\${client_resolve}</code>	Holds the result of the resolve call for <code>\${client_name}</code> : OK, FAIL, FORGED, or TEMP.
<code>\${deliveryMode}</code>	Specifies the current delivery mode <code>sendmail</code> is using instead of the value of the <code>DeliveryMode</code> option.
<code>\${dsn_notify}</code> , <code>\${dsn_envid}</code> , <code>\${dsn_ret}</code>	Holds the corresponding DSN parameter values.
<code>\${if_addr}</code>	Provides the interface's address for the incoming connection if the interface does not belong to the loopback net. This macro is especially useful for virtual hosting.
<code>\${if_addr_out}</code> , <code>\${if_name_out}</code> , <code>\${if_family_out}</code>	Avoids the reuse of <code>\${if_addr}</code> . Holds the following values respectively. The address of the interface for the outgoing connection The host name of the interface for the outgoing connection The family of the interface for the outgoing connection
<code>\${if_name}</code>	Provides the interface's host name for the incoming connection and is especially useful for virtual hosting.

TABLE 27–10 Defined Macros for `sendmail` (Continued)

Macro	Description
<code>\${load_avg}</code>	Checks and reports the current average number of jobs in the run queue.
<code>\${msg_size}</code>	Holds the value of the message size (<code>SIZE=parameter</code>) in an ESMTP dialogue before the message has been collected. Thereafter, the macro holds the message size as computed by <code>sendmail</code> and is used in <code>check_compat</code> . For information about <code>check_compat</code> , refer to Table 27–14.
<code>\${nrcpts}</code>	Holds the number of validated recipients.
<code>\${ntries}</code>	Holds the number of delivery attempts.
<code>\${rcpt_mailer}</code> , <code>\${rcpt_host}</code> , <code>\${rcpt_addr}</code> , <code>\${mail_mailer}</code> , <code>\${mail_host}</code> , <code>\${mail_addr}</code>	Holds the results of parsing the RCPT and MAIL arguments—that is, the resolved right-hand side (RHS) triplet from the mail delivery agent (<code> \$#mailer</code>), the host (<code> \$@host</code>), and the user (<code> \$:addr</code>).

New Macros Used to Build the `sendmail` Configuration File

In this section, you can find the following.

- Table 27–11
- “New MAX Macros” on page 408

TABLE 27–11 New Macros Used to Build the `sendmail` Configuration File

Macro	Description
<code>LOCAL_MAILER_EOL</code>	Overrides the default end-of-line string for the local mailer.
<code>LOCAL_MAILER_FLAGS</code>	Adds <code>Return-Path:</code> header by default.
<code>MAIL_SETTINGS_DIR</code>	Contains the path (including the trailing slash) for the mail settings directory.
<code>MODIFY_MAILER_FLAGS</code>	Improves the <code>*_MAILER_FLAGS</code> . This macro sets, adds, or deletes flags.
<code>RELAY_MAILER_FLAGS</code>	Defines additional flags for the relay mailer.

New MAX Macros

Use the following new macros to configure the maximum number of commands that can be received before `sendmail` slows its delivery. You can set these MAX macros at compile time. The maximum values in the following table also represent the current default values.

TABLE 27-12 New MAX Macros

Macro	Maximum Value	Commands Checked by Each Macro
MAXBADCOMMANDS	25	Unknown commands
MAXNOOPCOMMANDS	20	NOOP, VERB, ONEX, XUSR
MAXHELOCOMMANDS	3	HELO, EHLO
MAXVRFYCOMMANDS	6	VRFY, EXPN
MAXETRNCOMMANDS	8	ETRN

Note – You can disable a macro’s check by setting the macro’s value to zero.

New and Revised m4 Configuration Macros for `sendmail`

This section contains a table of new and revised m4 configuration macros for `sendmail`. Use the following syntax to declare these macros.

symbolic_name ('value')

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

TABLE 27-13 New and Revised m4 Configuration Macros for `sendmail`

m4 Macro	Description
FEATURE ()	For details, refer to “Changes to the FEATURE () Declaration” on page 409.
LOCAL_DOMAIN ()	This macro adds entries to class w (\$=w).
MASQUERADE_EXCEPTION ()	A new macro that defines hosts or subdomains that cannot be masqueraded.

TABLE 27–13 New and Revised m4 Configuration Macros for `sendmail` (Continued)

m4 Macro	Description
<code>SMART_HOST()</code>	This macro can now be used for bracketed addresses, such as <code>user@[host]</code> .
<code>VIRTUSER_DOMAIN()</code> or <code>VIRTUSER_DOMAIN_FILE()</code>	When these macros are used, include <code>\$={VirtHost}</code> in <code>\$=R</code> . As a reminder, <code>\$=R</code> is the set of host names that are allowed to relay.

Changes to the `FEATURE()` Declaration

Refer to the following tables for information about the specific changes to the `FEATURE()` declarations.

- Table 27–14
- Table 27–15

To use the new and revised `FEATURE` names, use the following syntax.

```
FEATURE('name', 'argument')
```

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323, in Chapter 25.

TABLE 27–14 New and Revised `FEATURE()` Declarations

Name of <code>FEATURE()</code>	Description
<code>compat_check</code>	Argument: Refer to the example in the following paragraph. This new <code>FEATURE()</code> enables you to look for a key in the access map that consists of the sender address and the recipient address and is delimited by the following string, <code><@></code> . <code>sender@sdomain<@>recipient@rdomain</code> is an example.
<code>delay_checks</code>	Argument: <code>friend</code> , which enables a spam-friend test, or <code>hater</code> , which enables spam-hater test. A new <code>FEATURE()</code> that delays all checks. By using <code>FEATURE('delay_checks')</code> , the rule sets <code>check_mail</code> and <code>check_relay</code> are not called when a client connects or issues a <code>MAIL</code> command, respectively. Instead, these rule sets are called by the <code>check_rcpt</code> rule set. For details, refer to the <code>/usr/lib/mail/README</code> file.

TABLE 27–14 New and Revised `FEATURE()` Declarations (Continued)

Name of <code>FEATURE()</code>	Description
<code>dnsbl</code>	Argument: This <code>FEATURE()</code> accepts a maximum of two arguments: ■ DNS server name ■ Rejection message A new <code>FEATURE()</code> that you can include multiple times to check the return values for DNS lookups. Note that this <code>FEATURE()</code> enables you to specify the behavior of temporary lookup failures.
<code>enhdnsbl</code>	Argument: domain name A new <code>FEATURE()</code> that is an enhanced version of <code>dnsbl</code>, which enables you to check the return values for DNS lookups. For more information, refer to <code>/usr/lib/mail/README</code>.
<code>generics_entire_domain</code>	Argument: None A new <code>FEATURE()</code> that you can also use to apply <code>genericstable</code> to subdomains of <code>\$=G</code>.
<code>ldap_routing</code>	Argument: For details, refer to the “Release Notes” in http://www.sendmail.org. A new <code>FEATURE()</code> that implements LDAP address routing.
<code>local_lmtp</code>	Argument: Path name of an LMTP-capable mailer. The default is <code>mail.local</code>, which is LMTP capable in this Solaris release. A <code>FEATURE()</code> that now sets the delivery status notification (DSN) diagnostic-code type for the local mailer to the proper value of SMTP.
<code>local_no_masquerade</code>	Argument: None A new <code>FEATURE()</code> that you can use to avoid masquerading for the local mailer.
<code>lookupdotdomain</code>	Argument: None A new <code>FEATURE()</code> that you can also use to look up the <code>.domain</code> in the access map.
<code>nocanonicalify</code>	Argument: <code>canonicalify_hosts</code> or nothing A <code>FEATURE()</code> that now includes the following features. Enables a list of domains, as specified by <code>CANONIFY_DOMAIN</code> or <code>CANONIFY_DOMAIN_FILE</code>, to be passed to the <code>\$[</code> and <code>\$]</code> operators for canonification. Enables addresses that have only a host name, such as <code><user@host></code>, to be canonified, if <code>canonicalify_hosts</code> is specified as its parameter. Adds a trailing dot to addresses with more than one component.

TABLE 27–14 New and Revised `FEATURE()` Declarations (Continued)

Name of <code>FEATURE()</code>	Description
<code>no_default_msa</code>	Argument: None A new <code>FEATURE()</code> that turns off <code>sendmail</code> 's default setting from <code>m4</code> -generated configuration files to "listen" on several different ports, an implementation of RFC 2476.
<code>nouucp</code>	Argument: <code>reject</code> , which does not allow the <code>!</code> token, or <code>nospecial</code> , which does allow the <code>!</code> token. A <code>FEATURE()</code> that determines whether to allow the <code>!</code> token in the local part of an address.
<code>nullclient</code>	Argument: None A <code>FEATURE()</code> that now provides the full rule sets of a normal configuration, allowing anti-spam checks to be performed.
<code>preserve_local_plus_detail</code>	Argument: None A new <code>FEATURE()</code> that enables you to preserve the <code>+detail</code> portion of the address when <code>sendmail</code> passes the address to the local delivery agent.
<code>preserve_luser_host</code>	Argument: None A new <code>FEATURE()</code> that enables you to preserve the name of the recipient host, if <code>LUSER_RELAY</code> is used.
<code>queuegroup</code>	Argument: None A new <code>FEATURE()</code> that enables you to select a queue group that is based on the full email address or on the domain of the recipient.
<code>relay_mail_from</code>	Argument: The <i>domain</i> is an optional argument. A new <code>FEATURE()</code> that allows relaying if the mail sender is listed as a <code>RELAY</code> in the access map and is tagged with the <code>From:</code> header line. If the optional <i>domain</i> argument is given, the domain portion of the mail sender is also checked.
<code>virtuser_entire_domain</code>	Argument: None A <code>FEATURE()</code> that you can now use to apply <code>\${VirtHost}</code> , a new class for matching <code>virtusertable</code> entries that can be populated by <code>VIRTUSER_DOMAIN</code> or <code>VIRTUSER_DOMAIN_FILE</code> . <code>FEATURE('virtuser_entire_domain')</code> can also apply the class <code>\${VirtHost}</code> to entire subdomains.

The following `FEATURE()` declarations are no longer supported.

TABLE 27–15 Unsupported FEATURE() Declarations

Name of FEATURE()	Replacement
rbl	FEATURE('dnsbl') and FEATURE('enhdnsbl') replace this FEATURE() that has been removed.
remote_mode	MASQUERADE_AS('\$S') replaces FEATURE('remote_mode') in /usr/lib/mail/cf/subsidiary.mc. \$S is the SMART_HOST value in sendmail.cf.
sun_reverse_alias_files	FEATURE('genericstable').
sun_reverse_alias_nis	FEATURE('genericstable').
sun_reverse_alias_nisplus	FEATURE('genericstable').

Changes to the MAILER() Declaration

The MAILER() declaration specifies support for delivery agents. To declare a delivery agent, use the following syntax.

```
MAILER('symbolic_name')
```

Note the following changes.

- In this new version of sendmail, the MAILER('smtp') declaration now includes an additional mailer, `dsmt`, which provides on-demand delivery by using the `F=%` mailer flag. The `dsmt` mailer definition uses the new `DSMT_MAILER_ARGS`, which defaults to `IPC $h`.
- Numbers for rule sets that are used by MAILERs have been removed. You now have no required order for listing your MAILERs except for MAILER('uucp'), which must follow MAILER('smtp') if `uucp-dom` and `uucp-uudom` are used.

For more information about mailers, refer to “Mailers” on page 356. If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

New Delivery Agent Flags

The following table describes new delivery agent flags, which by default are not set. These single-character flags are Boolean. You can set or unset a flag by including or excluding it in the `F=` statement of your configuration file, as shown in the following example.

```
Mlocal,    P=/usr/lib/mail.local, F=lsDFMAw5:|@qSXfmnz9, S=10/30, R=20/40,
Mprog,     P=/bin/sh, F=lsDFMoqeu9, S=10/30, R=20/40, D=$z:/,
Msmt,      P=[IPC], F=mDFMuX, S=11/31, R=21, E=\r\n, L=990,
```

```

Mesmtp,      P=[IPC], F=mDFMuXa, S=11/31, R=21, E=\r\n, L=990,
Msmtp8,      P=[IPC], F=mDFMuX8, S=11/31, R=21, E=\r\n, L=990,
Mrelay,      P=[IPC], F=mDFMuXa8, S=11/31, R=61, E=\r\n, L=2040,

```

TABLE 27-16 New Mailer Flags

Flag	Description
%	Mailers that use this flag do not attempt delivery to the initial recipient of a message or to queue runs unless the queued message is selected by using an ETRN request or one of the following queue options: -qI, -qR, or -qS.
1	This flag disables the ability of the mailer to send null characters (for example, \0).
2	This flag disables the use of ESMTP and requires that SMTP be used instead.
6	This flag enables mailers to strip headers to 7 bit.

New Equates for Delivery Agents

The following table describes new equates that you can use with the `M` delivery agent definition command. The following syntax shows you how to append new equates or new arguments to the equates that already exist in the configuration file.

```
Magent_name, equate, equate, ...
```

The following example includes the new `W=` `equate`, which specifies the maximum time to wait for the mailer to return after all data has been sent.

```
Msmtp, P=[IPC], F=mDFMuX, S=11/31, R=21, E=\r\n, L=990, W=2m
```

When you modify the definition of a value for `m4` configuration, use the syntax that is provided in the following example.

```
define('SMTP_MAILER_MAXMSG', '1000')
```

The preceding example places a limit of 1000 on the number of messages that are delivered per connection on an `smtp` mailer.

If you need to build a new `sendmail.cf` file, refer to “Building the `sendmail.cf` Configuration File (Task)” on page 323 in Chapter 25.

Note – Typically, you modify the `equate` definitions in the `mailer` directory only when you fine-tune.

TABLE 27–17 New Equates for Delivery Agents

Equate	Description
/=	Argument: Path to a directory Specifies a directory to apply <code>chroot()</code> to before the mailer program is executed
m=	Argument: Any of the following <code>m4</code> values that have previously been defined with the <code>define()</code> routine SMTP_MAILER_MAXMSGs, for the <code>smtp</code> mailer LOCAL_MAILER_MAXMSGs, for the <code>local</code> mailer RELAY_MAILER_MAXMSGs, for the <code>relay</code> mailer Limits the number of messages that are delivered per connection on an <code>smtp</code> , <code>local</code> , or <code>relay</code> mailer
W=	Argument: An increment of time Specifies the maximum time to wait for the return of the mailer after all data has been sent

New Queue Features

The following list provides details about new queue features.

- This release supports multiple queue directories. To use multiple queues, supply a `QueueDirectory` option value in the configuration file that ends with an asterisk (*), as is shown in the following example.

```
O QueueDirectory=/var/spool/mqueue/q*
```

The option value, `/var/spool/mqueue/q*`, uses all of the directories (or symbolic links to directories) that begin with “q” as queue directories. Do not change the queue directory structure while `sendmail` is running. Queue runs create a separate process for running each queue unless the verbose flag (`-v`) is used on a non-daemon queue run. The new items are randomly assigned to a queue.

- The new queue file-naming system uses file names that are guaranteed to be unique for 60 years. This system allows queue IDs to be assigned without complex file-system locking and simplifies the movement of queued items between queues.
- In version 8.12 of `sendmail`, only `root` can run the mail queue. For further details, refer to the changes that are described in the `mailq(1)` man page. For new task information, refer to “Administering the Queue Directories (Task Map)” on page 338.
- To accommodate envelope splitting, queue file names are now 15–characters long, rather than 14–characters long. File systems with a 14–character name limit are no longer supported.

For task information, refer to “Administering the Queue Directories (Task Map)” on page 338.

Changes for LDAP in sendmail

The following list describes changes in the use of the Lightweight Directory Access Protocol (LDAP) with sendmail.

- LDAPROUTE_EQUIVALENT() and LDAPROUTE_EQUIVALENT_FILE() permit you to specify equivalent host names, which are replaced by the masquerade domain name for LDAP routing lookups. For more information, refer to /usr/lib/mail/README.
- As noted in the Release Notes that are part of the sendmail distribution at ftp://ftp.sendmail.org, the LDAPX map has been renamed to LDAP. Use the following syntax for LDAP.

Kldap ldap options

- This release supports the return of multiple values for a single LDAP lookup. Place the values to be returned in a comma-separated string with the -v option, as is shown.

Kldap ldap -v"mail,more_mail"

- If no LDAP attributes are specified in an LDAP map declaration, all attributes that are found in the match are returned.
- This version of sendmail prevents commas in quoted key and value strings in the specifications of the LDAP alias file from breaking up a single entry into multiple entries.
- This version of sendmail has a new option for LDAP maps. The option, -Vseparator enables you to specify a separator, so that a lookup can return both an attribute and a value that are separated by the relevant separator.
- Instead of using the %s token to parse an LDAP filter specification, you can also use the new token, %0, to encode the key buffer. The %0 token applies a literal meaning to LDAP special characters.

The following example shows how these tokens differ for a lookup on “*.”

TABLE 27–18 Comparison of Tokens

LDAP Map Specification	Specification Equivalent	Result
-k"uid=%s"	-k"uid=*"	Matches any record with a user attribute
-k"uid=%0"	-k"uid=\2A"	Matches a user with the name “*.”

The following table describes new LDAP map flags.

TABLE 27–19 New LDAP Map Flags

Flag	Description
-1	Requires a single match to be returned. If more than one match is returned, the results are the equivalent of no records being found.
-r <i>never always search find</i>	Sets the LDAP alias dereference option.
-Z <i>size</i>	Limits the number of matches to return.

New Built-in Mailer Feature

The old [TCP] built-in mailer is not available. Use the P=[IPC] built-in mailer instead. The interprocess communications ([IPC]) built-in mailer now enables delivery to a UNIX domain socket on systems that support it. You can use this mailer with LMTP delivery agents that listen on a named socket. An example mailer might resemble the following.

```
Mexecmail, P=[IPC], F=lsDFMmqSXzA5@/:|, E=\r\n,
S=10, R=20/40, T=DNS/RFC822/X-Unix, A=FILE /var/run/lmtpd
```

The first mailer argument in the [IPC] mailer is now checked for a legitimate value. The following table provides possible values for the first mailer argument.

TABLE 27–20 Possible Values for the First Mailer Argument

Value	Description
A=FILE	Use for UNIX domain socket delivery
A=TCP	Use for TCP/IP connections
A=IPC	Is no longer available as a first mailer argument

New Rule Sets

The following table lists the new rule sets and describes what they do.

TABLE 27–21 New Rule Sets

Set	Description
check_eoh	Correlates information that is gathered between headers and checks for missing headers. This rule set is used with the macro storage map and is called after all of the headers have been collected.

TABLE 27-21 New Rule Sets (Continued)

Set	Description
check_etrn	Uses the ETRN command (as check_rcpt uses RCPT).
check_expn	Uses the EXPN command (as check_rcpt uses RCPT).
check_vrfy	Uses the VRFY command (as check_rcpt uses RCPT).

The following list describes new rule set features.

- Numbered rule sets are also named, but they can still be accessed by their numbers.
- The H header configuration file command allows for a default rule set to be specified for header checks. This rule set is called only if the individual header has not been assigned its own rule set.
- Comments in rule sets (that is, text within parentheses) are not removed if the configuration file version is nine or greater. For example, the following rule matches the input token (1), but does not match the input token.


```
R$+ (1)      $@ 1
```
- sendmail accepts the SMTP RSET command even when it rejects commands because of TCP wrappers or the check_relay rule set.
- You receive a warning if you set the OperatorChars option multiple times. Also, do not set OperatorChars after the rule sets are defined.
- The name of the rule set, as well as its lines, are ignored if an invalid rule set is declared. The rule set lines are not added to S0.

Changes to Files

Note the following changes.

- The helpfile is now located in /etc/mail/helpfile. The old name (/etc/mail/sendmail.hf) has a symbolic link that points to the new name.
- The trusted-users file is now located in /etc/mail/trusted-users. During an upgrade, if the old name (/etc/mail/sendmail.ct) is detected, but not the new name, a hard link from the old name to the new name is created. Otherwise, no change is made. The default content is root.
- The local-host-names file is now located in /etc/mail/local-host-names. During an upgrade, if the old name (/etc/mail/sendmail.cw) is detected, but not the new name, a hard link from the old name to the new name is created. Otherwise, no change is made. The default content is zero length.
- The new name for /usr/lib/mail/cf/main-v7sun.mc is /usr/lib/mail/cf/main.mc.

- The new name for `/usr/lib/mail/cf/subsidiary-v7sun.mc` is `/usr/lib/mail/cf/subsidiary.mc`.

IPv6 Addresses in Configuration

In version 8.12 of `sendmail`, IPv6 addresses that are used in configuration should be prefixed with the `IPv6:` tag to identify the address properly. If you are not identifying an IPv6 address, a prefix tag is not used.

Changes to `mail.local`

The following table describes the new command-line options for the `mail.local` program, which is used by `sendmail` as a delivery agent for local mail.

TABLE 27–22 New Command-Line Options for `mail.local`

Option	Description
-7	Prevents the Local Mail Transfer Protocol (LMTP) mode from advertising 8BITMIME support in the LHL0 response
-b	Causes a permanent error instead of a temporary error if a mailbox exceeds its quota

`mail.local` is the default for LMTP mode. However, for this release, if you choose to use `mail.local` as the local delivery agent without being in LMTP mode, you need to do one of the following to set the `S` flag.

Use the following syntax for the configuration file.

```
MODIFY_MAILER_FLAGS('LOCAL', '+S')      # for the configuration file
```

Alternately, perform the following two steps for `m4` configuration.

```
define('MODIFY_MAILER_FLAGS', 'S')dnl    # first step
MAILER(local)dnl                          # second step
```

Note – `MODIFY_MAILER_FLAGS` is a new macro that is used to build the configuration file. For details, refer to “New Macros Used to Build the `sendmail` Configuration File” on page 407.

For a complete review, refer to the `mail.local(1M)` man page.

Changes to mailstats

The mailstats program, which provides statistics on mailer usage, is packaged with the sendmail program. The following table describes new options in mailstats.

TABLE 27-23 New mailstats Options

Option	Description
-C <i>filename</i>	Specifies a sendmail configuration file
-p	Provides clear statistics in a program-readable mode
-P	Also provides clear statistics in a program-readable mode, but this option does not truncate the statistics file

For more information, refer to the mailstats(1) man page.

Changes to makemap

The makemap command creates keyed database files for sendmail. The following table describes new makemap options. When you declare options, use the following syntax.

`makemap options class filename`

When you use the preceding syntax, remember the following.

- *options* are preceded by a dash (for example, -dN).
- *class* refers to the type of database (for example, btree, dbm, or hash).
- *filename* refers to the full path (or relative name) for the database file.

TABLE 27-24 New makemap Options

Option	Description
-C	Uses the specified sendmail configuration file for finding the TrustedUser option
-c	Uses the specified hash and btree cache size
-e	Allows an empty value from the right-hand side (RHS)

TABLE 27-24 New makemap Options (Continued)

Option	Description
-l	Lists map types that are supported
-t	Specifies a different delimiter, instead of white space
-u	Dumps (unmaps) the contents of the database to standard output

Note – If makemap is running as root, the ownership of the generated maps is automatically changed to the `TrustedUser`, as specified in the `sendmail` configuration file. For more information about the `TrustedUser` option, refer to Table 27-2.

For more information, refer to the `makemap(1M)` man page.

New Command, `editmap`

Use the new maintenance command, `editmap`, to query and edit single records in keyed database maps for `sendmail`. From the command line, use the following syntax.

```
editmap options maptype mapname key "value"
```

- *options* are preceded by a dash (for example, `-Nf`). The man page provides a list of options and explains how each option functions.
- *maptype* refers to the type of database. `editmap` can use `btree`, `dbm`, and `hash`.
- *mapname* refers to the full path (or relative name) for the database file.
- *key* refers to a single string or multitoken string that you can use for searches.
- “*value*” refers to the string that appears to the right of the key in a keyed database file. In the following example, `man` is the key and `man@example.com` is the assigned value for that key.

```
man    man@host.com
```

For a detailed description and a list of options, refer to the `editmap(1M)` man page.

Other Changes and Features of Interest

The following list describes other changes and features of interest.

- As noted in RFC 2476, `sendmail` now listens for submissions on port 587.
- As was noted in the Release Notes that are part of the `sendmail` distribution at `ftp://ftp.sendmail.org`, the `XUSR SMTP` command is deprecated. Mail user agents should begin using RFC 2476 Message Submission for initial user message submission.
- The `Content-Length:` header is no longer provided in messages that are piped to programs with any version of the Sun configuration files. However, this header is still provided for appended messages and ordinary mailbox deliveries that use any version of the Sun configuration files.
- `sendmail` now accepts connections when disk space is low, but in such situations it allows only `ETRN` commands.
- Entries in the alias file can be continued by putting a backslash directly before the new line.
- The timeout for sending a message by way of SMTP has been changed to check for delivery progress every five minutes. This change detects an inability to send information more quickly and reduces the number of processes that are waiting to time out.
- You can now copy the contents of a class to another class by using the syntax of the following example.

```
C{Dest} $={Source}
```

In the preceding example, all items in class `$={Source}` are copied into class `$={Dest}`.

- The maps are no longer optional by default. Also, if a problem occurs with a map, you receive an error message.
- Canonification is no longer attempted for any host or domain in class `P` (`$=P`).
- The `=` equate is not included in an option expansion if no value is associated with the option.
- Route addresses are stripped. For example, `<@a,@b,@c:user@d>` is converted to `<user@d>`.

Composed September 5, 2002

Modem–Related Network Services Topics

Chapter 29	Provides overview information for PPP
Chapter 30	Provides planning information for PPP
Chapter 31	Provides step-by-step instructions for setting up a dial-up PPP link
Chapter 32	Provides step-by-step instructions for setting up a leased-line PPP link
Chapter 33	Provides step-by-step instructions for setting up authentication on a PPP link
Chapter 34	Provides step-by-step instructions for creating PPPoE tunnels to support PPP links over DSL equipment.
Chapter 35	Provides instructions for PPP link maintenance and problem solving
Chapter 36	Provides detailed reference information for working with PPP
Chapter 37	Provides step-by-step instructions for migrating from the earlier Solaris PPP (asppp) to Solaris PPP 4.0
Chapter 38	Provides background information on UUCP
Chapter 39	Provides step-by-step instructions for setting up and troubleshooting UUCP
Chapter 40	Provides reference material on UUCP database files, UUCP configuration files, UUCP shell scripts, and UUCP troubleshooting information

Composed September 5, 2002

Solaris PPP 4.0 (Overview)

Solaris PPP 4.0 enables two computers in different physical locations to communicate with each other by using Point-to-Point Protocol (PPP) over a variety of media. The Solaris 9 operating environment includes Solaris PPP 4.0 as part of the base installation.

This chapter introduces Solaris PPP 4.0. Topics that are discussed include the following:

- “Solaris PPP 4.0 Basics” on page 425
- “PPP Configurations and Terminology” on page 428
- “PPP Authentication” on page 434
- “Support for DSL Users Through PPPoE” on page 436

Solaris PPP 4.0 Basics

Solaris PPP 4.0 implements the Point-to-Point Protocol (PPP), a data link protocol, which is a member of the TCP/IP protocol suite. PPP describes how data is transmitted between two endpoint machines, over communications media such as telephone lines.

Since the early 1990s, PPP has been a widely used Internet standard for sending datagrams over a communications link. The PPP standard is described in RFC 1661 by the Point-to-Point Working Group of the Internet Engineering Task Force (IETF). PPP is commonly used when remote computers call an Internet service provider (ISP) or a corporate server that is configured to receive incoming calls.

Solaris PPP 4.0 is based on the publicly available Australian National University (ANU) PPP-2.4 and implements the PPP standard. Both asynchronous and synchronous PPP links are supported.

Solaris PPP 4.0 Compatibility

Various versions of standard PPP are available and in wide use throughout the Internet community. ANU PPP-2.4 is a popular choice for Linux, Tru64 UNIX, and all three major BSD variants:

- FreeBSD
- OpenBSD
- NetBSD

Solaris PPP 4.0 brings the highly configurable features of ANU PPP-2.4 to machines that run the Solaris operating environment. Machines that run Solaris PPP 4.0 can easily set up PPP links to any machine that runs an implementation of standard PPP.

Some non-ANU-based PPP implementations that successfully interoperate with Solaris PPP 4.0 include the following:

- Solaris PPP, also known as `asppp`, available with the Solaris 2.4 through Solaris 8 operating environments
- Solstice™ PPP 3.0.1
- Microsoft Windows 98 DUN
- Cisco IOS 12.0 (synchronous)

Which Version of Solaris PPP to Use

Solaris PPP 4.0 is the PPP implementation that is supported by the Solaris 9 operating environment. The Solaris 9 operating environment does not include the earlier asynchronous PPP (`asppp`) software. Asynchronous PPP configuration is discussed in the Solaris 8 System Administrator Collection at <http://www.docs.sun.com>.

Why Use Solaris PPP 4.0?

If you currently use `asppp`, consider migrating to Solaris PPP 4.0. Note the following differences between the two Solaris PPP technologies:

- **Transfer modes**
`asppp` supports asynchronous communications only. Solaris PPP 4.0 supports both asynchronous communications and synchronous communications.
- **Configuration process**
Setting up `asppp` requires configuring the `asppp.cf` configuration file, three UUCP files, and the `ifconfig` command. Moreover, you have to preconfigure interfaces for all users who might log in to a machine.
Setting up Solaris PPP 4.0 requires defining options for the PPP configuration files, or issuing the `pppd` command with options. You can also use a combination of both the configuration file and command-line methods. Solaris PPP dynamically creates and removes interfaces. You do not have to directly configure PPP

interfaces for each user.

- **Solaris PPP 4.0 features not available from asppp**
 - MS-CHAPv1 and MS-CHAPv2 authentication
 - PPP over Ethernet (PPPoE) , to support ADSL bridges
 - PAM authentication
 - Plug-in modules
 - IPv6 addressing
 - Data compression that uses Deflate or BSD compress

Solaris PPP 4.0 Upgrade Path

If you are converting an existing asppp configuration to Solaris PPP 4.0, you can use the translation script that is provided with this release. For complete instructions, refer to “How to Convert From asppp to Solaris PPP 4.0” on page 568.

Where to Go for More Information About PPP

Many resources with information about PPP can be found in print and online. The following subsections give some suggestions.

Professional Reference Books About PPP

For more information about widely used PPP implementations, including ANU PPP, refer to the following books:

- Carlson, James. *PPP Design, Implementation, and Debugging*. 2nd ed. Addison-Wesley, 2000.
- Sun, Andrew. *Using and Managing PPP*. O'Reilly & Associates, 1999.

Web Sites About PPP

Go to the following web sites for general information about PPP:

- For a list of frequently asked questions (FAQ) and other information about pppd, go to the following site that is provided by the Internet Engineering group at Sun Microsystems, <http://playground.sun.com/pppd>.
- For ANU PPP information, go to the PPP repository of Australian National University, <http://.pserver.samba.org/cgi-bin/cvsweb/ppp/>.
- For technical information, FAQs, discussions about Solaris system administration, and earlier versions of PPP, go to Sun Microsystem's system administrators' resource, <http://www.sun.com/bigadmin/home/index.html>.

- For modem configuration and advice on many different implementations of PPP, refer to Stokely Consulting's Web Project Management & Software Development web site at <http://www.stokely.com/unix.serial.port.resource/ppp.slip.htm>.

Requests for Comments (RFCs) About PPP

Some useful Internet RFCs about PPP include the following:

- 1661 and 1662, which describe the major features of the PPP protocol
- 1334, which describes authentication protocols, such as Password Authentication Protocol (PAP) and Challenge-Handshake Authentication Protocol (CHAP)
- 2516, an informational RFC that describes PPP over Ethernet (PPPoE)

To obtain copies of PPP RFCs, specify the number of the RFC on the IETF RFC web page at <http://www.ietf.org/rfc.html>.

Man Pages About PPP

For technical details about the Solaris PPP 4.0 implementation, refer to the following man pages:

- `pppd(1M)`
- `chat(1M)`
- `pppstats(1M)`
- `pppoec(1M)`
- `pppoed(1M)`
- `sppptun(1M)`
- `snoop(1M)`

You can find the PPP-related man pages at <http://www.docs.sun.com> or through the `man` command.

PPP Configurations and Terminology

This section introduces PPP configurations. The section also defines terms that are used in this guide.

Solaris PPP 4.0 supports a number of configurations.

- Switched access, or *dial-up*, configurations
- Hardwired, or *leased-line* configurations




FIGURE 29-1 Parts of the PPP Link

The previous figure shows a basic PPP link. The link has the following parts:

- Two machines, usually in separate physical locations, called *peers*. A peer could be a personal computer, engineering workstation, large server, or even a commercial router, depending on a site's requirements.
- Serial interface on each peer. On Solaris machines, this interface could be *cua*, *hih*, or other interface, depending on whether you configure asynchronous or synchronous PPP.
- Physical link, such as a serial cable, a modem connection, or a leased line from a network provider, such as a T1 or T3 line.

Dial-up PPP Overview

The most commonly used PPP configuration is the *dial-up link*. In a dial-up link, the local peer *dials up* the remote peer to establish the connection and run PPP. In the dial-up process, the local peer calls the remote peer's telephone number to initiate the link.

A common dial-up scenario includes a home computer that calls a peer at an ISP, configured to receive incoming calls. Another scenario is a corporate site where a local machine transmits data over a PPP link to a peer in another building.

In this guide, the local peer that initiates the dial-up connection is referred to as the *dial-out machine*. The peer that receives the incoming call is referred to as the *dial-in server*. This machine is actually the target peer of the dial-out machine and might or might not be a true server.

PPP is not a client-server protocol. Some PPP documents use the terms "client" and "server" to refer to telephone call establishment. A dial-in server is not a true server like a file server or name server. Dial-in server is a widely used PPP term because dial-in machines often "serve" network accessibility to more than one dial-out machine. Nevertheless, the dial-in server is the target peer of the dial-out machine.

Parts of the Dial-up PPP Link




FIGURE 29-2 Basic Analog Dial-up PPP Link

The configuration for Location 1, the dial-out side of the link, is composed of the following elements:

- Dial-out machine, typically a personal computer or workstation in an individual's home.
- Serial interface on the dial-out machine. `/dev/cua/a` or `/dev/cua/b` is the standard serial interface for outgoing calls on machines that run Solaris software.
- Asynchronous modem or ISDN terminal adapter (TA) that is connected to a telephone jack.
- Telephone lines and services of a telephone company.

The configuration for Location 2, the dial-in side of the link, is composed of the following elements:

- Telephone jack or similar connector, which is connected to the telephone network
- Asynchronous modem or ISDN TA
- Serial interface on the dial-in server, either `ttya` or `ttyb` for incoming calls
- Dial-in server, which is connected to a network, such as a corporate intranet, or, in the instance of an ISP, the global Internet

Using ISDN Terminal Adapters With a Dial-out Machine

External ISDN TAs have faster speeds than modems, but you configure TAs in basically the same way. The major difference in configuring an ISDN TA is in the chat script, which requires commands specific to the TA's manufacturer. Refer to "Chat Script for External ISDN TA" on page 542 for information on chat scripts for ISDN TAs.

What Happens During Dial-up Communications

PPP configuration files on both the dial-out and dial-in peers contain instructions for setting up the link. The following process occurs as the dial-up link is initiated:

1. User or process on the dial-out machine runs the `pppd` command to start the link.
2. Dial-out machine reads its PPP configuration files. The dial-out machine then sends instructions over the serial line to its modem, including the phone number of the dial-in server.
3. Modem dials the phone number to establish a telephone connection with the modem on the dial-in server.
If necessary, the dial-out machine sends commands to the dial-in server to invoke PPP on the server.
4. Modem attached to the dial-in server begins link negotiation with the modem on the dial-out machine.
The series of text strings that the dial-out machine sends to the modem and dial-in server are contained in a file called a *chat script*.
5. When modem-to-modem negotiation completes, the modem on the dial-out machine reports "CONNECT."
6. PPP on both peers enters *Establish* phase, where Link Control Protocol (LCP) negotiates basic link parameters and the use of authentication.
7. If necessary, the peers authenticate each other.
8. PPP's Network Control Protocols (NCPs) negotiate the use of network protocols, such as IPv4 or IPv6.

The dial-out machine can then run `telnet` or a similar command to a host that is known to the dial-in server.

Leased-Line PPP Overview

A hardwired, *leased-line* PPP configuration involves two peers that are connected by a link. This link consists of a switched or an unswitched digital service leased from a provider. Solaris PPP 4.0 works over any full-duplex, point-to-point leased-line medium. Typically, a company rents a hardwired link from a network provider to connect to an ISP or other remote site.

Comparison of Dial-up and Leased-Line Links

Both dial-up and leased-line links involve two peers that are connected by a communications medium. The next table summarizes the differences between the link types.

Leased Line	Dial-up Line
Always connected, unless a system administrator or power failure takes the leased-line down	Initiated on demand, when a user tries to call a remote peer
Uses synchronous communications	Uses asynchronous communications
Rented from a provider	Uses existing telephone lines
Requires synchronous units	Uses less costly modems
Requires specialized interfaces	Uses standard serial interfaces that are included on most computers

Parts of a Leased-Line PPP Link




FIGURE 29-3 Basic Leased-Line Configuration

The leased-line link contains the following parts:

- **Two peers**, each peer at one end of the link. Each peer might be a workstation or server. Often the peer functions as a router between its network or the Internet, and the opposite peer.
- **Synchronous interface on each peer**. Some machines that run Solaris software require you to purchase a synchronous interface card, such as HSI/S, to connect to a leased line. Other machines, such as UltraSPARC™ workstations, have built-in synchronous interfaces.

- **CSU/DSU synchronous digital unit on each peer**, which connects the synchronous port to the leased line.

A CSU might be built-in to the DSU, or owned by you, or leased from a provider, depending on your locale. The DSU gives the Solaris machine a standard synchronous serial interface. With Frame Relay, the Frame Relay Access Device (FRAD) performs the serial interface adaptation.

- **Leased line**, providing switched or unswitched digital services. Some examples are SONET/SDH, Frame Relay PVC, and T1.

Note – SONET is called an *octet synchronous* link. PPP uses a framing mechanism that is similar to asynchronous framing over a SONET line. PPP does not use the expected bit-synchronous protocol.

What Happens During Leased-Line Communications

On most types of leased lines, peers do not actually dial each other. Rather, a company purchases a leased-line service to explicitly connect between two fixed locations. Sometimes the two peers at either end of the leased line are at different physical locations of the same company. Another scenario is a company that sets up a router on a leased line that is connected to an ISP.

Leased lines are less commonly used than dial-up links, though the hardwired links are easier to set up. Hardwired links do not require chat scripts. Authentication is often not used because both peers are known to each other when a line is leased. After the two peers initiate PPP over the link, the link stays active. A leased-line link remains active unless the line fails, or either peer explicitly terminates the link.

A peer on a leased line that runs Solaris PPP 4.0 uses most of the same configuration files that define a dial-up link.

The following process occurs to initiate communication over the leased line:

1. Each peer machine runs the `pppd` command as part of the booting process or another administrative script.
2. The peers read their PPP configuration files.
3. The peers negotiate communications parameters.
4. An IP link is established.

PPP Authentication

Authentication is the process of verifying that a user is who he or she claims to be. The UNIX login sequence is a simple form of authentication:

1. The `login` command prompts the user for a name and password.
2. `login` then attempts to authenticate the user by looking up the typed user name and password in the password database.
3. If the database contains the user name and password, then the user is *authenticated* and given access to the system. If the database does not contain the user name and password, the user is denied access to the system.

By default, Solaris PPP 4.0 does not demand authentication on machines that do not have a default route specified. Thus, a local machine without a default route does not authenticate remote callers. Conversely, if a machine does have a default route defined, by default the machine does authenticate remote callers.

You might use PPP authentication protocols to verify the identity of callers who are trying to set up a PPP link to your machine. Conversely, you must configure PPP authentication information if your local machine must call peers that authenticate callers.

Authenticators and Authenticatees

The calling machine on a PPP link is considered the *authenticatee* because the caller must prove its identity to the remote peer. The peer is considered the *authenticator*. The authenticator looks up the caller's identity in the appropriate PPP files for the security protocol and authenticates or does not authenticate the caller.

You typically configure PPP authentication for a dial-up link. When the call begins, the dial-out machine is the authenticatee. The dial-in server is the authenticator. The server has a database in the form of a *secrets* file. This file lists all users who are granted permission to set up a PPP link to the server. Think of these users as *trusted callers*.

Some dial-out machines require remote peers to provide authentication information when responding to the dial-out machine's call. Then, their roles are reversed: the remote peer becomes the authenticatee and the dial-out machine the authenticator.

Note – PPP 4.0 does not prevent authentication by leased-line peers, but authentication is not often used in leased-line links. The nature of leased-line contracts usually means that both participants on the ends of the line are known to each other. Both participants often are trusted. However, because PPP authentication is not that difficult to administer, you should seriously consider implementing authentication for leased lines.

PPP Authentication Protocols

The PPP authentication protocols are Password Authentication Protocol (PAP) and Challenge-Handshake Authentication Protocol (CHAP). Each protocol uses a *secrets* database that contains identification information, or *security credentials*, for each caller that is permitted to link to the local machine. For a detailed explanation of PAP, see "Password Authentication Protocol (PAP)" on page 546. For a CHAP explanation, see "Challenge-Handshake Authentication Protocol (CHAP)" on page 549.

Why Use PPP Authentication?

Providing authentication on a PPP link is optional. Moreover, though authentication does verify that a peer is to be trusted, PPP authentication does not provide confidentiality of data. For confidentiality, use encryption software, such as IPsec, PGP, SSL, and the Solaris secure shell.

Note – Solaris PPP 4.0 does not implement the PPP Encryption Control Protocol (ECP), which is described in RFC 1968.

Consider implementing PPP authentication in the following situations:

- Your company accepts incoming calls from users over the public, switched telephone network.
- Your corporate security policy requires remote users to provide authentication credentials when accessing your network through a corporate firewall or when engaging in secure transactions.
- You want to authenticate callers against a standard UNIX password database, such as `/etc/passwd`, NIS, NIS+, LDAP, or PAM. Use PAP authentication for this scenario.
- Your company's dial-in servers also provide the network's Internet connection. Use PAP authentication for this scenario.
- The serial line is less secure than the password database on the machine or networks at either end of the link. Use CHAP authentication for this scenario.

Support for DSL Users Through PPPoE

Many network providers and individuals who are working at home use Digital Subscriber Line (DSL) technology to provide fast network access. To support DSL users, Solaris PPP 4.0 includes the PPP over Ethernet (PPPoE) feature. PPPoE technology enables multiple hosts to run PPP sessions over one Ethernet link to one or more destinations.

If one of the following factors applies to your situation, you should use PPPoE:

- You support DSL users, possibly including yourself. Your DSL service provider might require users to configure a PPPoE tunnel to receive services over the DSL line.
- Your site is an ISP that intends to offer PPPoE to customers.

This section introduces terms that are associated with PPPoE and an overview of a basic PPPoE topology.

PPPoE Overview

PPPoE is a proprietary protocol from RedBack Networks. PPPoE is a discovery protocol, rather than another version of standard PPP. In a PPPoE scenario, a machine that initiates PPP communications first must locate, or *discover*, a peer that runs PPPoE. The PPPoE protocol uses Ethernet broadcast packets to locate the peer.

After the discovery process, PPPoE sets up an Ethernet-based tunnel from the initiating host, or *PPPoE client*, to the peer, the *PPPoE access server*. *Tunneling* is the practice of running one protocol on top of another protocol. Using PPPoE, Solaris PPP 4.0 tunnels PPP over Ethernet IEEE 802.2, both of which are data link protocols. The resulting PPP connection behaves like a dedicated link between the PPPoE client and the access server. For detailed information about PPPoE, see “Creating PPPoE Tunnels for DSL Support” on page 554.

Parts of a PPPoE Configuration

Three participants are involved in a PPPoE configuration: a consumer, a telephone company, and a service provider, as the following figure shows.




FIGURE 29-4 Participants in a PPPoE Tunnel

PPPoE Consumers

As system administrator, you might assist consumers with their PPPoE configurations. One common type of PPPoE consumer is an individual who needs to run PPPoE over a DSL line. Another PPPoE consumer is a company that purchases a DSL line through which employees can run PPPoE tunnels, as illustrated in the previous figure.

The main reason for a corporate consumer to use PPPoE is to offer PPP communications through a high-speed DSL device to a number of hosts. Often, a single PPPoE client has an individual *DSL modem*. Or, a group of clients on a hub might share a DSL modem that is also connected to the hub by an Ethernet line.

Note – DSL devices are technically bridges, not modems. However, because common practice is to refer to these devices as modems, this guide uses the term “DSL modem.”

PPPoE runs PPP over a tunnel on the Ethernet line that is connected to the DSL modem. That line is connected to a splitter, which, in turn connects to a telephone line.

PPPoE at a Telephone Company

The telephone company is the middle layer of the PPPoE scenario. The telephone company splits the signal that is received over the phone line by using a device that is called a *Digital Subscriber Line Access Multiplexer (DSLAM)*. The DSLAM breaks out the signals onto separate wires, analog wires for telephone service, and digital wires for PPPoE. From the DSLAM, the digital wires extend the tunnel over an ATM data network to the ISP.

PPPoE at a Service Provider

The ISP receives the PPPoE transmission from the ATM data network over a bridge. At the ISP, an access server that runs PPPoE functions as the peer for the PPP link. The access server is very similar in function to the dial-in server that was introduced in Figure 29–2, but the access server does not use modems. The access server converts the individual PPPoE sessions into regular IP traffic, for example Internet access.

If you are a system administrator for an ISP, you might be responsible for configuring and maintaining an access server.

Security on a PPPoE Tunnel

The PPPoE tunnel is inherently insecure. You can use PAP or CHAP to provide user authentication for the PPP link that is running over the tunnel.

Composed September 5, 2002

Planning for the PPP Link (Tasks)

Setting up a PPP link involves a set of discrete tasks, which includes planning tasks and other activities that are not related to PPP. This chapter explains how to plan for the most common PPP links, for authentication, and for PPPoE.

The task chapters that follow Chapter 30 use sample configurations to illustrate how to set up a particular link. These sample configurations are introduced in this chapter.

Topics that are covered include the following:

- “Planning a Dial-up PPP Link” on page 442
- “Planning a Leased-Line Link” on page 446
- “Planning for Authentication on a Link” on page 449
- “Planning for DSL Support Over a PPPoE Tunnel” on page 454

Overall PPP Planning (Task Map)

PPP requires planning tasks before you actually can set up the link. Moreover, if you want to use a PPPoE tunneling, you first have to set up the PPP link and then provide tunneling. The following task map lists the large planning tasks that are discussed in this chapter. You might need to use only the general task for the link type to be configured. Or you might require the task for the link, authentication, and perhaps PPPoE.

TABLE 30-1 Task Map for PPP Planning

Task	Description	For Instructions
Plan for a dial-up PPP link	Gather information that is required to set up a dial-out machine or a dial-in server	“Planning a Dial-up PPP Link” on page 442

TABLE 30–1 Task Map for PPP Planning *(Continued)*

Task	Description	For Instructions
Plan for a leased-line link	Gather information that is required to set up a client on a leased line	“Planning a Leased-Line Link” on page 446
Plan for authentication on the PPP link	Gather information that is required to configure PAP or CHAP authentication on the PPP link	“Planning for Authentication on a Link” on page 449
Plan for a PPPoE tunnel	Gather information that is required to set up a PPPoE tunnel over which a PPP link can run	“Planning for DSL Support Over a PPPoE Tunnel” on page 454

Planning a Dial-up PPP Link

Dial-up links are the most commonly used PPP links. This section includes the following information:

- Planning information for a dial-up link
- Explanation of the sample link to be used in Chapter 31

Typically, you only configure the machine at one end of the dial-up PPP link, the dial-out machine, or the dial-in server. For an introduction to dial-up PPP, refer to “Dial-up PPP Overview” on page 429.

Before You Set Up the Dial-out Machine

Before you configure a dial-out machine, gather the information that is listed in the following table.

Note – The planning information in this section does not include information to be gathered about authentication or PPPoE. For details on authentication planning, refer to “Planning for Authentication on a Link” on page 449. For PPPoE planning, refer to “Planning for DSL Support Over a PPPoE Tunnel” on page 454.

TABLE 30-2 Information for a Dial-out Machine

Information	Action
Maximum modem speed	Refer to documentation that was provided by the modem manufacturer.
Modem connection commands (AT commands)	Refer to documentation that was provided by the modem manufacturer.
Name to use for dial-in server at the other end of the link	Create any name that helps you identify the dial-in server.
Login sequence that was required by dial-in server	Contact the dial-in server's administrator or ISP documentation if dial-in server is at the ISP.

Before You Set Up the Dial-in Server

Before you configure a dial-in server, gather the information that is listed in the following table.

Note – The planning information in this section does not include information to be gathered about authentication or PPPoE. For details on authentication planning, refer to “Planning for Authentication on a Link” on page 449. For PPPoE planning, refer to “Planning for DSL Support Over a PPPoE Tunnel” on page 454.

TABLE 30-3 Information for a Dial-in Server

Information	Action
Maximum modem speed	Refer to documentation that was provided by the modem manufacturer.
User names of people who are permitted to call the dial-in server	Obtain the names of the prospective users before you set up their home directories, as discussed in “How to Configure Users of the Dial-in Server” on page 470.
Dedicated IP address for PPP communications	Obtain an address from the individual at your company who is responsible for delegating IP addresses.

Example—Configuration for Dial-up PPP

The tasks to be introduced in Chapter 31 execute a small company's requirement to let employees work at home a few days a week. Some employees require the Solaris operating environment on their home machines. These workers also need to log in remotely to their work machines on the corporate intranet.

Composed September 5, 2002

The tasks set up a basic dial-up link with the following features:

- The *dial-out* machines are at the houses of employees who need to call the corporate intranet.
- The *dial-in* server is a machine on the corporate intranet that is configured to receive incoming calls from employees.
- UNIX-style login is used to authenticate the dial-out machine. Stronger Solaris PPP 4.0 authentication methods are not required by the company's security policy.

The next figure shows the link that is set up in Chapter 31.




FIGURE 30-1 Sample Dial-up Link

In this figure, a remote host dials out through its modem over telephone lines to Big Company's intranet. Another host is configured to dial out to Big Company but currently is inactive. The calls from remote users are answered in the order received by the modem that is attached to the dial-in server at Big Company. A PPP connection is established between the peers. The dial-out machine can then remotely log in to a host machine on the intranet.

Where to Go For More Information About Dial-up PPP

Task	For Information
Set up a dial-out machine	Table 31-2
Set up a dial-in machine	Table 31-4
Get an overview of dial-up links	"Dial-up PPP Overview" on page 429
Get detailed information about PPP files and commands	"Using PPP Options in Files and on the Command Line" on page 523

Planning a Leased-Line Link

Setting up a leased-line link involves configuring the peer at one end of a switched or unswitched service leased from a provider.

This section includes the following information:

- Planning information for a leased-line link
- Explanation of the sample link that is shown in Figure 30-2

For an introduction to leased-line links, refer to "Leased-Line PPP Overview" on page 432. For tasks on setting up the leased line, see Chapter 32.

Before You Set Up the Leased-Line Link

When your company rents a leased-line link from a network provider, you typically configure only the system at your end of the link. The peer at the other end of the link is maintained by another administrator. This individual might be a system administrator at a remote location in your company or a system administrator at an ISP.

Hardware That Is Needed for a Leased-Line Link

In addition to the link media, your end of the link requires the following hardware:

- Synchronous interface for your system
- Synchronous unit (CSU/DSU)

■ Your system

Some network providers include a router, synchronous interface, and a CSU/DSU as part of the customer premises equipment (CPE). However, necessary equipment varies, based on the provider and any governmental restrictions in your locale. The network provider can give you information about the unit that is needed, if this equipment is not provided with the leased line.

Information to Be Gathered for the Leased-Line Link

Before you configure the local peer, you might need to gather the items that are listed in the next table.

TABLE 30-4 Planning for a Leased-Line Link

Information	Action
Device name of the interface	Refer to the Interface card documentation.
Configuration instructions for the synchronous interface card	Refer to the Interface card documentation. You need this information to configure the HSI/S interface. You might not need to configure other types of interface cards.
(Optional) IP address of the remote peer	Refer to the service provider documentation. Alternatively, contact the system administrator of the remote peer. This information is needed only if the IP address is not negotiated between the two peers.
(Optional) Name of the remote peer	Refer to the service provider documentation. Alternatively, you can contact the system administrator of the remote peer.
(Optional) Speed of the link	Refer to the service provider documentation. Alternatively, you can contact the system administrator of the remote peer.
(Optional) Compression that is used by the remote peer	Refer to the service provider documentation. Alternatively, you can contact the system administrator of the remote peer.

Example—Configuration for a Leased-Line Link

The tasks in Chapter 32 show how to implement the goal of a medium-sized organization (LocalCorp) to provide Internet access for its employees. Currently, the employees' computers are connected on a private corporate intranet.

LocalCorp requires speedy transactions and access to the many resources on the Internet. The organization signs a contract with Far ISP, a service provider, which allows LocalCorp to set up its own leased line to Far ISP. Then, LocalCorp leases a T1 line from Phone East, a telephone company. Phone East puts in the leased line between LocalCorp and Far ISP. Then, Phone East provides a CSU/DSU that is already configured to Local Corp.

The tasks set up a leased-line link with the following characteristics.

- LocalCorp has set up a system as a gateway router, which forwards packets over the leased line to hosts on the Internet.
- Far ISP also has set up a peer as a router to which leased lines from customers are attached.




FIGURE 30-2 Sample Leased-Line Configuration

In the figure, a router is set up for PPP at LocalCorp. The router connects to the corporate intranet through its hme0 interface. The second connection is through the machine's HSI/S interface (hih1) to the CSU/DSU digital unit. The CSU/DSU then connects to the installed leased line. The administrator at LocalCorp configures the HSI/S interface and PPP files. The administrator then types `/etc/init.d/pppd` to initiate the link between LocalCorp and FarISP.

Where to Get More Information About Leased Lines

Task	For Information
Set up a client on a leased line	Chapter 32
Get an overview of leased lines	"Leased-Line PPP Overview" on page 432

Planning for Authentication on a Link

This section contains planning information for providing authentication on the PPP link. Chapter 33 contains tasks for implementing PPP authentication at your site.

PPP offers two types of authentication, PAP, which is described in detail in "Password Authentication Protocol (PAP)" on page 546 and CHAP, which is described in "Challenge-Handshake Authentication Protocol (CHAP)" on page 549.

Before you set up authentication on a link, you must choose which authentication protocol best meets your site's security policy. Then, you set up the secrets file and PPP configuration files for the dial-in machines, or callers' dial-out machines, or both types of machines. For information on choosing the appropriate authentication protocol for your site, see "Why Use PPP Authentication?" on page 436.

This section includes the following information:

- Planning information for both PAP and CHAP authentication
- Explanations of the sample authentication scenarios that are shown in Figure 30–3 and Figure 30–4

For tasks on setting up authentication, see Chapter 33.

Before You Set Up PPP Authentication

Setting up authentication at your site should be an integral part of your overall PPP strategy. Before implementing authentication, you should assemble the hardware, configure the software, and test the link.

TABLE 30-5 Prerequisites Before Configuring Authentication

Information	For Instructions
Tasks for configuring a dial-up link	Chapter 31.
Tasks for testing the link	Chapter 35.
Security requirements for your site	Your corporate security policy. If you do not have a policy, setting up PPP authentication gives you an opportunity to create a security policy.
Suggestions about whether to use PAP or CHAP at your site	"Why Use PPP Authentication?" on page 436. For more detailed information about these protocols, refer to "Authenticating Callers on a Link" on page 545.

Example—PPP Authentication Configurations

This section contains the sample authentication scenarios to be used in the procedures in Chapter 33.

Example—Configuration Using PAP Authentication

The tasks in "Configuring PAP Authentication" on page 482 show how to set up PAP authentication over the PPP link. The procedures use as an example a PAP scenario that was created for the fictitious "Big Company" in "Example—Configuration for Dial-up PPP" on page 443.

Big Company wants to enable its users to work from home. The system administrators want a secure solution for the serial lines to the dial-in server. UNIX-style login that uses the NIS password databases has served Big Company's network well in the past. The system administrators want a UNIX-like authentication scheme for calls that come in to the network over the PPP link. So, the administrators implement the following scenario that uses PAP authentication.




FIGURE 30-3 Example—PAP Authentication Scenario (Working From Home)

The system administrators create a dedicated dial-in DMZ that is separated from the rest of the corporate network by a router. The term DMZ comes from the military term “demilitarized zone.” The DMZ is an isolated network that is set up for security purposes. The DMZ typically contains resources that a company offers to the public, such as web servers, anonymous FTP servers, databases, and modem servers. Network designers often place the DMZ between a firewall and a company’s Internet connection.

The only occupants of the DMZ that is pictured in Figure 30-3 are the dial-in server **myserver** and the router. The dial-in server requires callers to provide PAP credentials, including user names and passwords, when setting up the link.

Furthermore, the dial-in server uses the `login` option of PAP. Therefore, the callers' PAP user names and passwords must correspond exactly to their UNIX user names and passwords in the dial-in server's password database.

After the PPP link is established, the caller's packets are forwarded to the router. The router forwards the transmission to its destination on the corporate network or on the Internet.

Example—Configuration Using CHAP Authentication

The tasks in "Configuring CHAP Authentication" on page 490 show how to set up CHAP authentication. The procedures use as an example a CHAP scenario to be created for the fictitious LocalCorp that was introduced in "Example—Configuration for a Leased-Line Link" on page 447.

LocalCorp provides connectivity to the Internet over a leased line to an ISP. The Technical Support department within LocalCorp generates heavy network traffic. Therefore, Technical Support requires its own, isolated private network. The department's field technicians travel extensively and need to access the Technical Support network from remote locations for problem-solving information. To protect sensitive information in the private network's database, remote callers must be authenticated in order to be granted permission to log in.

Therefore, the system administrators implement the following CHAP authentication scenario for a dial-up PPP configuration.




FIGURE 30-4 Example—CHAP Authentication Scenario (Calling a Private Network)

The only link from the Technical Support network to the outside world is the serial line to the dial-in server's end of the link. The system administrators configure the laptop computer of each field service representative for PPP with CHAP security, including a CHAP secret. The chap-secrets database on the dial-in server contains the CHAP credentials for all machines that are allowed to call in to the Technical Support network.

Where to Get More Information About Authentication

Task	For Instructions
Set up PAP authentication	"Configuring PAP Authentication" on page 482
Set up CHAP authentication	"Configuring CHAP Authentication" on page 490
Learn details about PPP authentication	"Authenticating Callers on a Link" on page 545 and the pppd(1M) man page

Planning for DSL Support Over a PPPoE Tunnel

Some DSL providers require you to set up PPPoE tunneling for your site in order to run PPP over the providers' DSL lines and high-speed digital networks. For an overview of PPPoE, see "Support for DSL Users Through PPPoE" on page 436.

A PPPoE tunnel involves three participants: a consumer, a telephone company, and an ISP. You either configure PPPoE for consumers—PPPoE clients at your company or consumers in their homes—or on a server at an ISP.

This section contains planning information for running PPPoE on both clients and access servers. The following topics are covered:

- Planning information for the PPPoE host and access server
- Explanation of the PPPoE scenario that is introduced in "Example—Configuration for a PPPoE Tunnel" on page 456

For tasks on setting up a PPPoE tunnel, see Chapter 34.

Before You Set Up a PPPoE Tunnel

Your preconfiguration activities depend on whether you configure the client side or server side of the tunnel. In either instance, you or your organization must contract with a telephone company. The telephone company provides the DSL lines for clients, and some form of bridging and possibly an ATM pipe for access servers. In most contracts, the telephone company assembles its equipment at your site.

Before Configuring a PPPoE Client

PPPoE client implementations usually consist of the following equipment:

- Personal computer or other system used by an individual
- DSL modem, which is usually installed by the telephone company or Internet access provider
- (Optional) A hub, if more than one client is involved, as is true for corporate DSL consumers
- (Optional) A splitter, usually installed by the provider

Many different DSL configurations are possible, which depend on the user or corporation's needs and the services that are offered by the provider.

TABLE 30-6 Planning for PPPoE Clients

Information	Action
If setting up a home PPPoE client for an individual or yourself, get any setup information that is outside the scope of PPPoE.	Ask the telephone company or ISP for any required setup procedures.
If setting up PPPoE clients at a corporate site, get the names of users to get PPPoE client systems. If you configure remote PPPoE clients, you might be responsible for giving users information about adding home DSL equipment.	Ask management at your company for a list of authorized users.
Find out what interfaces are available on the PPPoE client.	Run the <code>ifconfig -a</code> command on each machine for interface names.
(Optional) Get the password for the PPPoE client.	Ask users for passwords their preferred passwords. Or, assign passwords to the users. Note that this password is used for link authentication, not for UNIX login.

Before Configuring a PPPoE Server

Planning for a PPPoE access server involves working with the telephone company that provides your connection to its data service network. The telephone company installs its lines, often ATM pipes, at your site, and provides some sort of bridging into your access server. You need to configure the Ethernet interfaces that access the services that your company provides. For example, you need to configure interfaces for Internet access, as well as the Ethernet interfaces from the telephone company's bridge.

TABLE 30-7 Planning for a PPPoE Access Server

Information	Action
Interfaces that are used for lines from data service network	Run the <code>ifconfig -a</code> command to identify interfaces.
Types of services to provide from the PPPoE server	Ask management, network planners for their requirements and suggestions.
(Optional) Types of services to provide to the consumers	Ask management, network planners for their requirements and suggestions.
(Optional) Host names and passwords for remote clients	Ask network planners and other individuals at your site who are responsible for contract negotiations. The host names and passwords are used for PAP or CHAP authentication, not UNIX login.

Example—Configuration for a PPPoE Tunnel

This section contains a sample PPPoE tunnel, which is used as an illustration for the tasks in Chapter 34. Though the illustration shows all participants in the tunnel, you only administer one end, either the client side or server side.




FIGURE 30-5 Example—PPPoE Tunnel

In the sample, MiddleCo wants to provide its employees with high-speed Internet access. MiddleCo buys a DSL package from Phone East, which, in turn, contracts with service provider Far ISP. Far ISP offers Internet and other IP services to customers who buy DSL from Phone East.

Example—PPPoE Client Configuration

MiddleCo buys a package from Phone East that provides one DSL line for the site. The package includes a dedicated, authenticated connection to the ISP for MiddleCo's PPPoE clients. The system administrator cables the prospective PPPoE clients to a hub. Technicians from Phone East cable the hub to their DSL equipment.

Example—PPPoE Server Configuration

To implement the business arrangement FarISP has with Phone East, the system administrator at FarISP configures the access server `dslserve`. This server has the following four interfaces:

- `le0` – Primary network interface that connects to the local network
- `hme0` – Interface through which FarISP provides Internet service for its customers
- `hme1` – Interface contracted by MiddleCo for authenticated PPPoE tunnels
- `hme2` – Interface contracted by other customers for their PPPoE tunnels

Where to Get More Information About PPPoE

Task	For Instructions
Set up a PPPoE client	"Setting Up the PPPoE Client" on page 498
Set up a PPPoE access server	"Setting Up a PPPoE Access Server" on page 501
Get detailed information about PPPoE	"Creating PPPoE Tunnels for DSL Support" on page 554 and the <code>pppoed(1M)</code> , <code>pppoec(1M)</code> , and <code>sppptun(1M)</code> man pages

Composed September 5, 2002

Setting Up a Dial-up PPP Link (Tasks)

This chapter explains the tasks for configuring the most common PPP link, the dial-up link. Major topics include the following:

- “Configuring the Dial-out Machine” on page 460
- “Configuring the Dial-in Server” on page 467
- “Calling the Dial-in Server” on page 473

Major Tasks for Setting Up the Dial-up PPP Link (Task Map)

You set up the dial-up PPP link by configuring modems, modifying network database files, and modifying the PPP configuration files that are described in Table 36–1.

The next table lists the major tasks to configure both sides of a dial-up PPP link. Typically, you configure only one end of the link, either the dial-out machine or dial-in server.

TABLE 31–1 Task Map for Setting Up the Dial-up PPP Link

Task	Description	For Instructions
1. Gather preconfiguration information	Gather data that is needed prior to setting up the link, such as peer host names, target phone numbers, and modem speed.	“Planning a Dial-up PPP Link” on page 442
2. Configure the dial-out machine	Set up PPP on the machine that makes the call over the link.	Table 31–2

TABLE 31-1 Task Map for Setting Up the Dial-up PPP Link *(Continued)*

Task	Description	For Instructions
3. Configure the dial-in server	Set up PPP on the machine that receives incoming calls.	Table 31-4
4. Call the dial-in server	Type the <code>pppd</code> command to initiate communications.	"How to Call the Dial-in Server" on page 473

Configuring the Dial-out Machine

The tasks in this section explain how to configure a dial-out machine. The tasks use as an example the dial-in-from-home scenario that was introduced in Figure 30-1. You can perform the tasks at your company before passing on the machine to a prospective user. Alternatively, you can instruct experienced users in the setup of their home machines. Anyone setting up a dial-out machine must have root permission for that machine.

Tasks for Configuring the Dial-out Machine (Task Map)

TABLE 31-2 Task Map for Setting Up the Dial-out Machine

Task	Description	For Instructions
1. Gather preconfiguration information	Gather data that is needed prior to setting up the link, such as peer host names, target phone numbers, and modem speed.	"Planning a Dial-up PPP Link" on page 442
2. Configure the modem and serial port	Set up the modem and serial port.	"How to Configure the Modem and Serial Port (Dial-out Machine)" on page 462
3. Configure the serial-line communication	Configure the characteristics of the transmission across the serial line.	"How to Define Communications Over the Serial Line" on page 463
4. Define the conversation between the dial-out machine and the peer	Gather communications data for use when you create the chat script.	"How to Create the Instructions for Calling a Peer" on page 464

TABLE 31–2 Task Map for Setting Up the Dial-out Machine *(Continued)*

Task	Description	For Instructions
5. Configure information about a particular peer	Configure PPP options to call an individual dial-in server.	“How to Define the Connection With an Individual Peer” on page 465
6. Call the peer	Type the <code>pppd</code> command to initiate communications.	“How to Call the Dial-in Server” on page 473.

Dial-up PPP Template Files

Solaris PPP 4.0 provides template files. Each template contains common options for a particular PPP configuration file. The next table lists the sample templates that can be used for setting up a dial-up link, and their equivalent Solaris PPP 4.0 files.

Template File	PPP Configuration File	For Instructions
<code>/etc/ppp/options.tpl</code>	<code>/etc/ppp/options</code>	“ <code>/etc/ppp/options.tpl</code> Template” on page 528
<code>/etc/ppp/options.ttya.tpl</code>	<code>/etc/ppp/options.ttyname</code>	“ <code>options.ttya.tpl</code> Template File” on page 530
<code>/etc/ppp/myisp-chat.tpl</code>	File with the name of your choice to contain the chat script	“ <code>/etc/ppp/myisp-chat.tpl</code> Chat Script Template” on page 538
<code>/etc/ppp/peers/myisp.tpl</code>	<code>/etc/ppp/peers/peer-name</code>	“ <code>/etc/ppp/peers/myisp.tpl</code> Template File” on page 534

If you decide to use one of the template files, be sure to rename the template to its equivalent PPP configuration file. The one exception is the chat file template `/etc/ppp/myisp-chat.tpl`. You can give chat scripts any names that you want.

Configuring Devices on the Dial-out Machine

The first task for setting up a dial-out PPP machine is to configure the devices on the serial line: the modem and serial port.

Note – Tasks that apply to a modem usually apply to an ISDN TA.

Before performing the next procedure, you must have done the following.

- Installed the Solaris 9 operating environment on the dial-out machine

- Determined the optimum modem speed
- Decided which serial port to use on the dial-out machine
- Obtained the root password for the dial-out machine

For planning information, see Table 30–2.

▼ How to Configure the Modem and Serial Port (Dial-out Machine)

1. Program the modem.

Even though a variety of modem types is available, most modems are shipped with the correct settings for Solaris PPP 4.0. The following table lists basic settings for modems that use Solaris PPP 4.0.

TABLE 31–3 Modem Settings for Dial-up PPP

Parameter	Setting
DCD	Follow carrier
DTR	Low so that the modem hangs up-puts the modem on-hook
Flow Control	RTS/CTS for full-duplex hardware flow control
Attention Sequences	Disable

If you have problems in setting up the link and suspect that the modem is at fault, first consult the modem manufacturer's documentation. Also, a number of sites on the World Wide Web offer help with modem programming. Finally, you can find some suggestions for clearing modem problems in "How to Diagnose Modem Problems" on page 514.

2. **Attach the modem cables to the serial port on the dial-out machine and to the telephone jack.**
3. **Become superuser on the dial-out machine.**
4. **Run `admintool`, as explained in "Setting Up Terminals and Modems with Serial Ports Tool" in *System Administration Guide: Advanced Administration*.**
 - a. **Click the port where you have attached the modem, either port a or port b.**

The Modify Serial Port window is displayed.
 - b. **Specify modem direction as dial-out only.**

You can set up the modem as bidirectional, the default template for `admintool`. However, the dial-out-only choice is more secure against possible intruders.

Note – You can set the baud rate and timeout from `admintool`. However, the `pppd` daemon ignores these settings.

5. Click Okay to convey the changes.

Configuring Communications on the Dial-out Machine

The procedures in this section show how to configure communications over the serial line of the dial-out machine. Before you can use these procedures, you must have configured the modem and serial port, as described in “How to Configure the Modem and Serial Port (Dial-out Machine)” on page 462.

The next tasks show how to enable the dial-out machine to successfully initiate communications with the dial-in server. Communications are initiated as defined in the options in the PPP configuration files. You need to create the following files:

- `/etc/ppp/options`
- `/etc/ppp/options.ttyname`
- Chat script
- `/etc/ppp/peers/peer-name`

Solaris PPP 4.0 provides templates for the PPP configuration files, which you can customize to suit your needs. Refer to “Dial-up PPP Template Files” on page 461 for detailed information about these files.

▼ How to Define Communications Over the Serial Line

1. Become superuser on the dial-out machine.
2. Create a file that is called `/etc/ppp/options` with the following entry:

`lock`

The `/etc/ppp/options` file is used for defining global parameters that apply to all communications by the local machine. The `lock` option enables UUCP-style locking of the form `/var/spool/locks/LK.xxx.yyy.zzz`.

Note – If the dial-out machine does not have an `/etc/ppp/options` file, only the superuser can run the `pppd` command. However, the `/etc/ppp/options` can be empty.

For a complete description of `/etc/ppp/options`, refer to “`/etc/ppp/options` Configuration File” on page 528.

3. (Optional) Create a file that is called `/etc/ppp/options.ttyname` for defining how communications should be initiated from a specific serial port.

The next example shows an `/etc/ppp/options.ttyname` file for the port with the device name `/dev/cua/a`.

```
# vi /etc/ppp/options.cua.a
crtscts
```

The PPP option `crtscts` tells the `pppd` daemon to turn on hardware flow control for serial port `a`.

For more information about the `/etc/ppp/options.ttyname` file, go to “`/etc/ppp/options.ttyname` Configuration File” on page 529.

4. Set the modem speed, as described in “How to Set the Modem Speed” on page 469.

▼ How to Create the Instructions for Calling a Peer

Before the dial-out machine can initiate a PPP link, you must collect information about the dial-in server to become the peer. Then, you use this information to create the chat script, which describes the actual conversation between the dial-out machine and the peer.

1. Determine the speed at which the dial-out machine’s modem needs to run.

For more information, see “Configuring the Modem Speed” on page 535.

2. Obtain the following information from the dial-in server’s site:

- Server’s telephone number
- Authentication protocol that is used, if appropriate
- Login sequence that is required by the peer for the chat script

3. Obtain the names and IP addresses of name servers at the dial-in server’s site.

4. Put instructions for initiating calls to the particular peer in a chat script.

For example, you might create the following chat script, `/etc/ppp/mychat`, to call the dial-in server `myserver`.

```
SAY "Calling the peer\n"
TIMEOUT 10
ABORT BUSY
```


Composed September 5, 2002

```
ABORT 'NO CARRIER'
ABORT ERROR
REPORT CONNECT
"" AT&F1&M5S2=255
TIMEOUT 60
OK ATDT1-123-555-1234
CONNECT \c
SAY "Connected; logging in.\n"
TIMEOUT 5
ogin:--ogin: pppuser
TIMEOUT 20
ABORT 'ogin incorrect'
ssword: \qmypassword
"% " \c
SAY "Logged in. Starting PPP on peer system.\n"
ABORT 'not found'
"" "exec pppd"
~ \c
```

The script contains instructions for calling a Solaris dial-in server that requires a login sequence. For a description of each instruction, refer to “Basic Chat Script Enhanced for a UNIX-Style Login” on page 540. For complete details on creating a chat script, read the section “Defining the Conversation on the Dial-up Link” on page 536.

Note – You do not invoke the chat script directly. Rather, you use the file name of the chat script as an argument to the `connect` option, which invokes the script.

If a peer runs Solaris or a similar operating system, consider using the previous chat script as a template for your dial-out machines.

▼ How to Define the Connection With an Individual Peer

1. Become superuser on the dial-out machine.
2. Update DNS databases by creating the following `/etc/resolv.conf` file:

```
domain bigcompany.com
nameserver 10.10.111.15
nameserver 10.10.130.8
```

<code>domain bigcompany.com</code>	Specifies that the peer's DNS domain is <code>bigcompany.com</code> .
<code>nameserver 10.10.111.15</code>	Lists the IP addresses of name servers at <code>bigcompany.com</code> .
<code>nameserver 10.10.130.8</code>	

For complete details on DNS implementation, refer to “DNS Administration (Reference)” in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

3. Edit the `/etc/nsswitch.conf` file to have the DNS database searched first for host information.

```
hosts:      dns [NOTFOUND=return] files
```

4. Create the `/etc/ppp/peers` directory, and then add a file for the peer.

For example, you would create the following file to define the dial-in server `myserver`:

```
# cd /etc/ppp
# mkdir peers
# cd peers
# vi myserver
/dev/cua/a
57600
noipdefault
defaultroute
idle 120
noauth
connect "chat -U 'mypassword' -T 1-123-555-1213 -f /etc/ppp/mychat"
```

<code>/dev/cua/a</code>	Specifies that the device <code>/dev/cua/a</code> should be used as the serial interface for calls to <code>myserver</code> .
<code>57600</code>	Defines the speed of the link.
<code>noipdefault</code>	Specifies that for transactions with peer <code>myserver</code> , the dial-out machine initially has an IP address of 0.0.0.0. <code>myserver</code> assigns an IP address to the dial-out machine for every dial-up session.
<code>idle 120</code>	Indicates that the link must time out after an idle period of 120 seconds.
<code>noauth</code>	Specifies that the peer <code>myserver</code> does not need to provide authentication credentials when negotiating the connection with the dial-out machine.
<code>connect "chat -U 'mypassword' -T 1-123-555-1213 -f /etc/ppp/mychat"</code>	Specifies the <code>connect</code> option and its arguments, including the phone number of the peer, and the chat script <code>/etc/ppp/mychat</code> with calling instructions.

Where to Go From Here

Task	For Instructions
Configure another dial-out machine	"How to Configure the Modem and Serial Port (Dial-out Machine)" on page 462.
Test modem connectivity by dialing out to another computer	<code>cu(1C)</code> and <code>tip(1)</code> man pages. These utilities can help you test if your modem is properly configured. Also, use these utilities to test if you can establish a connection with another machine.
Get detailed information about the PPP configuration files	"Using PPP Options in Files and on the Command Line" on page 523.
Begin configuring the dial-in server	"Configuring Devices on the Dial-in Server" on page 468.

Configuring the Dial-in Server

The tasks in this section are for configuring the dial-in server. The dial-in server is a peer machine that receives the call over the PPP link from the dial-out machine. The tasks show how to configure the dial-in server `myserver` that was introduced in Figure 30–1.

Tasks for Configuring the Dial-in Server (Task Map)

TABLE 31–4 Task Map for Setting Up the Dial-in Server

Task	Description	For Instructions
1. Gather preconfiguration information	Gather data that is needed prior to setting up the link, such as peer host names, target phone numbers, and modem speed.	"Planning a Dial-up PPP Link" on page 442
2. Configure the modem and serial port	Set up the modem and serial port.	"How to Configure the Modem and Serial Port (Dial-in Server)" on page 468

TABLE 31-4 Task Map for Setting Up the Dial-in Server (Continued)

Task	Description	For Instructions
3. Configure calling peer information	Set up the user environments and PPP options for every dial-out machine that is permitted to call the dial-in server.	"How to Configure Users of the Dial-in Server" on page 470
4. Configure the serial-line communication.	Configure the characteristics of the transmission across the serial line.	"How to Define Communications Over the Serial Line (Dial-in Server)" on page 471

Configuring Devices on the Dial-in Server

The following procedure explains how to configure the modem and serial port on the dial-in server.

Before you do the next procedure, you must have completed the following activities on the peer dial-in server:

- Installed the Solaris 9 operating environment
- Determined the optimum modem speed
- Decided which serial port to use

How to Configure the Modem and Serial Port (Dial-in Server)

1. **Program the modem, as instructed in the modem manufacturer's documentation.**

For other suggestions, refer to "How to Configure the Modem and Serial Port (Dial-out Machine)" on page 462.

2. **Attach the modem to the serial port on the dial-in server.**

3. **Become superuser on the dial-in server.**

4. **Configure the serial port by using `admintool`, as described in "Setting Up Terminals and Modems with Serial Ports Tool" in *System Administration Guide: Advanced Administration*.**

Use `admintool` to do the following:

- a. **Select the serial port where you have attached the modem, either port a or port b.**
The Modify Serial Port window is displayed.
- b. **Specify modem direction as dial-in only.**

Note – Solaris PPP 4.0 does support bidirectional communications for a modem.

- c. Click Okay to convey the changes.

▼ How to Set the Modem Speed

The next procedure explains how to set the modem speed for a dial-in server. For suggestions on speeds to use with Sun Microsystems' computers, see "Configuring the Modem Speed" on page 535.

1. **Log in to the dial-in server.**
2. **Use the `tip` command to reach the modem.**
Instructions for using `tip` to set the modem speed are in the `tip(1)` man page.
3. **Configure the modem for a fixed DTE rate.**
4. **Lock the serial port to that rate, using `ttymon` or `admintool`, as discussed in "Setting Up Terminals and Modems with Serial Ports Tool" in *System Administration Guide: Advanced Administration*.**

Where to Go From Here

Task	For Instructions
Configure another serial port and modem on the dial-in server	"How to Configure the Modem and Serial Port (Dial-in Server)" on page 468
Configure information about users that call the dial-in server	"How to Configure Users of the Dial-in Server" on page 470

Setting Up Users of the Dial-in Server

Part of the process of setting up a dial-in server involves configuring information about each known remote caller.

Before starting the procedures in this section, you must have done the following:

- Obtained the UNIX user names for all users who are permitted to log in from remote dial-out machines.
- Set up the modem and serial line, as described in "How to Configure the Modem and Serial Port (Dial-in Server)" on page 468.

- Dedicated an IP address to be assigned to incoming calls from remote users. Consider creating a dedicated incoming IP address if the number of potential callers exceeds the number of modems and serial ports on the dial-in server. For complete information about creating dedicated IP addresses, go to “Creating an IP Addressing Scheme for Callers” on page 552.

▼ How to Configure Users of the Dial-in Server

1. Become superuser on the dial-in server.

2. Create a new account on the dial-in server for each remote PPP user.

You can use `admintool` or the Solaris Management Console to create a new user. For instructions on creating a new user through Solaris Management Console, see “Setting Up User Accounts (Task Map)” in *System Administration Guide: Basic Administration*. For instructions on creating a new user through `admintool`, see `admintool(1M)`.

Note – The remaining steps show how to create an account by using `admintool`. You can use the same parameters for creating an account with Solaris Management Console.

3. Use the Add User template to create the new user.

For example, the next table shows how you might fill out PPP-related parameters for an account that is called `pppuser` for `user1` on the dial-out machine `myhome`.

Template Parameter	Value	Definition
User Name	<code>pppuser</code>	The user account name for the remote user. This account name should correspond to the account name that is given in the login sequence of the chat script. For example, <code>pppuser</code> is the account name that is found in the chat script in “How to Create the Instructions for Calling a Peer” on page 464.
Login Shell	<code>/usr/bin/pppd</code>	The default login shell for the remote user. The login shell <code>/usr/bin/pppd</code> initially restricts the caller to a dedicated PPP environment.
Create Home Dir Path	<code>/export/home/pppuser</code>	The home directory <code>/export/home/pppuser</code> is set when the caller successfully logs in to the dial-in server.

4. Create for each caller a `$HOME/.ppprc` file that contains various options that are specific to the user’s PPP session.

For example, you might create the following `.ppprc` file for `pppuser`.

```
# cd /export/home/pppuser
# vi .ppprc
noccp
noccp turns off compression control on the link.
```

Where to Go From Here

Task	For Instructions
Set up more users of the dial-in server	"How to Configure Users of the Dial-in Server" on page 470
Configure communications over the dial-in server	"How to Define Communications Over the Serial Line (Dial-in Server)" on page 471

Configuring Communications Over the Dial-in Server

The next task shows how to enable the dial-in server to open communications with any dial-out machine. The options that are defined in the following PPP configuration files determine how communications are established.

- `/etc/ppp/options`
- `/etc/ppp/options.ttyname`

For detailed information about these files, refer to "Using PPP Options in Files and on the Command Line" on page 523.

Before you proceed, you should have done the following:

- Configured the serial port and modem on the dial-in server, as described in "How to Configure the Modem and Serial Port (Dial-in Server)" on page 468.
- Configured information about the prospective users of the dial-in server, as described in "How to Configure Users of the Dial-in Server" on page 470.

How to Define Communications Over the Serial Line (Dial-in Server)

1. **Become superuser on the dial-in server.**
2. **Create the `/etc/ppp/options` file with the following entry.**

```
nodefaultroute
```

`nodfaultroute` indicates that no route is defined for the server.

Note – If the dial-in server does not have an `/etc/ppp/options` file, only the superuser can run the `pppd` command. However, the `/etc/ppp/options` file can be empty.

3. **Create the file `/etc/options.ttyname` to define how calls that are received over serial port `ttyname` should be handled.**

The following `/etc/options.ttya` file defines how the dial-in server’s serial port `/dev/ttya` should handle incoming calls.

<code>:10.0.0.80</code>	
<code>xonxoff</code>	
<code>:10.0.0.80</code>	Assigns the IP address 10.0.0.80 to all peers that are calling in over serial port <code>ttya</code>
<code>xonxoff</code>	Allows the serial line to handle communications from modems with software flow control enabled

Where to Go From Here

If you have followed all the procedures in this chapter, you have completed the configuration of the dial-up link.

Task	For Instructions
Test modem connectivity by dialing out to another computer	<code>cu(1C)</code> and <code>tip(1)</code> man pages. These utilities can help you test if your modem is properly configured. Also, use these utilities to test if you can establish a connection with another machine.
Configure more options for the dial-in server	“Configuring the Dial-in Server” on page 467
Configure more dial-out machines	“Configuring the Dial-out Machine” on page 460
Have the remote machine call the dial-in server	“Calling the Dial-in Server” on page 473

Calling the Dial-in Server

You establish a dial-up PPP link by having the dial-out machine call the dial-in server. You can instruct the dial-out machine to call the server by specifying the demand option in the local PPP configuration files. However, the most common method for establishing the link is for the user to run the `pppd` command on the dial-out machine.

Before you proceed to the next task, you should have done either or both of the following:

- Set up the dial-out machine, as described in “Configuring the Dial-out Machine” on page 460
- Set up the dial-in server, as described in “Configuring the Dial-in Server” on page 467

▼ How to Call the Dial-in Server

1. Log in to the dial-out machine by using your regular user account, not `root`.
2. Call the dial-in server by running the `pppd` command.

For example, the following command initiates a link between the dial-out machine and dial-in server `myserver`:

```
% pppd 57600 call myserver
```

<code>pppd</code>	Starts the call by invoking the <code>pppd</code> daemon
<code>57600</code>	Sets the speed of the line between host and modem
<code>call myserver</code>	Invokes the <code>call</code> option of <code>pppd</code> . <code>pppd</code> then reads options in the file <code>/etc/ppp/peers/myserver</code> , which was created in “How to Define the Connection With an Individual Peer” on page 465

3. Contact a host on the server’s network, for example, the host `lindyhop` that is shown in Figure 30-1:

```
ping lindyhop
```

If the link is working correctly, the standard Telnet login sequence should be displayed in the terminal window. If the link is not working correctly, refer to Chapter 35.

4. Terminate the PPP session:

```
% pkill -TERM -x pppd
```

Where to Go From Here

If you have followed all the procedures in this chapter, you have completed the configuration of the dial-up link.

Task	For Instructions
Have users start working on their dial-out machines	"How to Call the Dial-in Server" on page 473
Fix problems on the link	Chapter 35
Learn more about the files and options that are used in this chapter	"Using PPP Options in Files and on the Command Line" on page 523

CHAPTER **32**

Setting Up a Leased-Line PPP Link (Tasks)

This chapter explains how to configure a PPP link that uses a leased line between peers. Major sections include the following:

- “Configuring Synchronous Devices on the Leased Line” on page 476
- “Configuring a Machine on the Leased Line” on page 477

Setting Up a Leased Line (Task Map)

Leased-line links are relatively easy to set up, in comparison with dial-up links. In most instances, you do not have to configure the CSU/DSU, dialing services, or authentication. If you do need to configure the CSU/DSU, refer to the manufacturer’s documentation for aid with this complex task.

The task map in the next table describes all the tasks that are involved in setting up the basic leased-line link.

Note – Some types of leased lines do require the CSU/DSU to “dial” the address of the opposite peer. Some examples are Frame Relay that uses Switched Virtual Circuits (SVCs) or Switched 56 service.

TABLE 32–1 Task Map for Setting Up the Leased-Line Link

Task	Description	For Instructions
1. Gather preconfiguration information	Gather data that is needed prior to setting up the link.	Table 30–4

TABLE 32-1 Task Map for Setting Up the Leased-Line Link *(Continued)*

Task	Description	For Instructions
2. Set up the leased-line hardware	Assemble the CSU/DSU and synchronous interface card.	"How to Configure Synchronous Devices" on page 476
3. Configure the interface card, if required	Configure the interface script to be used when the leased line is brought up.	"How to Configure Synchronous Devices" on page 476
4. Configure information about the remote peer	Define how communications between your local machine and the remote peer should work.	"How to Configure a Machine on a Leased Line" on page 478
5. Start up the leased line	Configure your machine to start up PPP over the leased line as part of the booting process.	"How to Configure a Machine on a Leased Line" on page 478

Configuring Synchronous Devices on the Leased Line

The task in this section involves configuring equipment that is required by the leased-line topology that is introduced in "Example—Configuration for a Leased-Line Link" on page 447. The synchronous devices that are required to connect to the leased line include the interface and modem.

Prerequisites for Synchronous Devices Setup

Before you perform the next procedure, you must have the following items:

- Working leased line that is installed at your site by the provider
- Synchronous unit (CSU/DSU)
- Solaris 9 operating environment release installed on your system
- Synchronous interface card of the type that is required by your system

▼ How to Configure Synchronous Devices

1. **Physically install the interface card into the local machine, if necessary.**
Follow the instructions in the manufacturer's documentation.
2. **Connect the cables from the CSU/DSU to the interface. If necessary, connect cables from the CSU/DSU to the leased-line jack or similar connector.**

3. Configure the CSU/DSU, as instructed in the documentation from the manufacturer or network provider.

Note – The provider from whom you rented the leased line might supply and configure the CSU/DSU for your link.

4. Configure the interface card, if necessary, as instructed in the interface documentation.

The configuration of the interface card involves the creation of a startup script for the interface. The router at LocalCorp in the leased-line configuration that is shown in Figure 30–2 uses an HSI/S interface card.

The following script `hsi-conf` starts up the HSI/S interface.

```
#!/bin/ksh
/opt/SUNWconn/bin/hsi_init hih1 speed=1536000 mode=fdx loopback=no \
nrzi=no txc=txc rxc=rxr txd=txd rxd=rxr signal=no 2>&1 > /dev/null
```

hih1	Indicates that HSI/S is the synchronous port used

speed=1536000	Sets the speed of the CSU/DSU to 1536000

Where to Go From Here

Task	For Instructions
Configure the local machine on the leased line	"How to Configure a Machine on a Leased Line" on page 478

Configuring a Machine on the Leased Line

The task in this section explains how to set up a router to function as the local peer on your end of a leased line. The task uses the leased line that was introduced in "Example—Configuration for a Leased-Line Link" on page 447 as an example.

Prerequisites for Configuring the Local Machine on a Leased Line

Before you perform the next procedure, you must have completed the following:

- Set up and configured the synchronous devices for the link, as described in “Configuring Synchronous Devices on the Leased Line” on page 476
- Obtained the root password for the local machine on the leased line
- Set up the local machine to run as a router on the network(s) to use the services of the leased-line provider

▼ How to Configure a Machine on a Leased Line

1. **Become superuser on the local machine (router).**
2. **Add an entry for the remote peer in the router's `/etc/hosts` file.**

```
# vi /etc/hosts
#
# Internet host table
#
127.0.0.1      localhost
192.168.130.10 local2-peer      loghost
192.168.130.11 local1-net
10.0.0.25     farISP
```

The sample `/etc/hosts` file is for the local router at the fictitious LocalCorp. Note the IP address and host name for the remote peer `farISP` at the service provider.

3. **Create the file `/etc/ppp/peers/peer-name` to hold information about the provider's peer.**

For the sample leased-line link, you create the file `/etc/ppp/peers/farISP`.

```
#vi /etc/ppp/peers/farISP
init '/etc/ppp/conf_hsi'
local
/dev/hih1
sync
noauth
192.168.130.10:10.0.0.25
nodefaultroute
passive
persist
noccp
nopcomp
novj
noaccomp
```

The following table explains the options and parameters that are used in `/etc/ppp/peers/farISP`.

Option	Definition
init '/etc/ppp/conf_hsi'	Starts up the link. init then configures the HSI interface by using the parameters in the script /etc/ppp/conf_hsi.
local	Tells the pppd daemon not to change the state of the Data Terminal Ready (DTR) signal. Also tells pppd to ignore the Data Carrier Detect (DCD) input signal.
/dev/hih1	Gives the device name of synchronous interface.
sync	Establishes synchronous encoding for the link.
noauth	Disables authentication on the link.
192.168.130.10:10.0.0.25	Defines the IP addresses of the local peer and the remote peer, separated by a colon.
passive	Tells the pppd daemon on the local machine to go quiet after issuing maximum number of LCP Configure-Requests and wait for the peer to start.
persist	Tells the pppd daemon to try to restart the link after a connection ends.
noccp, nopcomp, novj, noaccomp	Disables the Compression Control Protocol (CCP), Protocol Field compression, Van Jacobson compression, and address and control field compression, respectively. These forms of compression accelerate transmissions on a dial-up link but could slow down a leased line.

4. Create an initialization script that is called demand, which creates the PPP link as part of the booting process.

```
# cd /etc/ppp/
# vi demand
if [ -f /var/run/ppp-demand.pid ] &&
    /usr/bin/kill -s 0 `/bin/cat /var/run/ppp-demand.pid`
then
    :
else
    /usr/bin/pppd call farISP
fi
```

The demand script contains the pppd command for establishing a leased-line link. The following table explains the contents of \$PPPPDIR/demand.

Code Sample	Explanation
echo "Starting Solaris PPP 4.0\c"	Displays "Starting Solaris PPP 4.0" during the booting process.

Code Sample	Explanation
<pre>if ps -e grep '\<pppd\ > /dev/null 2>&1 ; then echo "\npppd daemon is still running" echo "or in the process of exiting" exit 0 echo "\nEstablishing PPP session...\n" /usr/bin/pppd call farISP</pre>	Search for an already existing pppd daemon. If pppd is found, then send out a message and exit the demand script. Display “Establishing PPP session” during booting. Run the pppd command by using the options that are in /etc/ppp/peers/farISP.

The Solaris PPP 4.0 startup script /etc/rc2.d/S47pppd invokes the demand script as part of the Solaris booting process. The following lines in /etc/rc2.d/S47pppd search for the presence of a file that is called \$PPPDIR/demand.

```
if [ -f $PPPDIR/demand ]; then
    . $PPPDIR/demand
fi
```

If found, \$PPPDIR/demand is executed. During the course of executing \$PPPDIR/demand, the link is established.

Where to Go From Here

If you have followed all the procedures in this chapter, you have completed the configuration of the leased-line link.

Task	For Instructions
Instruct users to start communicating with machines on the Internet or other network that is served by the remote peer	Have users run telnet, ftp, rsh, or similar commands to reach machines outside the local network.
Fix problems on the link	“Fixing Leased-Line Problems” on page 518 for troubleshooting information.
Learn more about the files and options that are used in this chapter	“Using PPP Options in Files and on the Command Line” on page 523.

Setting Up PPP Authentication (Tasks)

This chapter contains tasks for setting up PPP authentication. Subjects that are covered include the following:

- “Configuring PAP Authentication” on page 482
- “Configuring CHAP Authentication” on page 490

The procedures show how to implement authentication over a dial-up link because dial-up links are more likely to be configured for authentication than leased-line links. You can configure authentication over leased lines if authentication is required by your corporate security policy. For leased-line authentication, use the tasks in this chapter as guidelines.

If you want to use PPP authentication but are not sure which protocol to use, review the section “Why Use PPP Authentication?” on page 436. More detailed information about PPP authentication is in the `pppd(1M)` man page and in “Authenticating Callers on a Link” on page 545.

Configuring PPP Authentication (Task Map)

This section contains task maps to help you quickly access procedures for PPP authentication.

TABLE 33–1 Task Map for General PPP Authentication

Task	For Instructions
Configure PAP authentication	“Setting Up PAP Authentication (Task Maps)” on page 482
Configure CHAP authentication	“Setting Up CHAP Authentication (Task Maps)” on page 490

Configuring PAP Authentication

The tasks in this section explain how to implement authentication on a PPP link by using the Password Authentication Protocol (PAP). The tasks use the example that is shown in “Example—PPP Authentication Configurations” on page 450 to illustrate a working PAP scenario for a dial-up link. Use the instructions as the basis for implementing PAP authentication at your site.

Before you perform the next procedures, you must have done the following:

- Set up and tested the dial-up link between the dial-in server and dial-out machines that belong to trusted callers
- Ideally, for dial-in server authentication, obtained superuser permission for the machine where the network password database is administered, for example, in LDAP, NIS, or local files
- Obtained superuser authority for the local machine, either dial-in server or dial-out machine

Setting Up PAP Authentication (Task Maps)

Use the next task maps to quickly access PAP-related tasks for the dial-in server and trusted callers on dial-out machines.

TABLE 33–2 Task Map for PAP Authentication (Dial-in Server)

Task	Description	For Instructions
1. Gather preconfiguration information	Collect user names and other data that is needed for authentication.	“Planning for Authentication on a Link” on page 449
2. Update the password database, if necessary	Ensure that all potential callers are in the server’s password database.	“How to Create a PAP Credentials Database (Dial-in Server)” on page 483

TABLE 33–2 Task Map for PAP Authentication (Dial-in Server) *(Continued)*

Task	Description	For Instructions
3. Create the PAP database	Create security credentials for all prospective callers in <code>/etc/ppp/pap-secrets</code> .	“How to Create a PAP Credentials Database (Dial-in Server)” on page 483
4. Modify the PPP configuration files	Add options specific to PAP to the <code>/etc/ppp/options</code> and <code>/etc/ppp/peers/peer-name</code> files.	“How to Add PAP Support to the PPP Configuration Files (Dial-in Server)” on page 485

TABLE 33–3 Task Map for PAP Authentication (Dial-out Machine)

Task	Description	For Instructions
1. Gather preconfiguration information	Collect user names and other data that is needed for authentication.	“Planning for Authentication on a Link” on page 449
2. Create the PAP database for the trusted caller’s machine	Create the security credentials for the trusted caller and, if necessary, security credentials for other users who call the dial-out machine, in <code>/etc/ppp/pap-secrets</code> .	“How to Configure PAP Authentication Credentials for the Trusted Callers” on page 487
3. Modify the PPP configuration files	Add options specific to PAP to the <code>/etc/ppp/options</code> and <code>/etc/ppp/peers/peer-name</code> files.	“How to Add PAP Support to the PPP Configuration Files (Dial-out Machine)” on page 488

Configuring PAP Authentication on the Dial-in Server

To set up PAP authentication, you must do the following:

- Create a PAP credentials database
- Modify PPP configuration files for PAP support

▼ How to Create a PAP Credentials Database (Dial-in Server)

This procedure modifies the `/etc/ppp/pap-secrets` file, which contains the PAP security credentials that are used to authenticate callers on the link. `/etc/ppp/pap-secrets` must exist on both machines on a PPP link.

The sample PAP configuration that was introduced in Figure 30–3 uses the `login` option of PAP. If you plan to use this option, you might also need to update your network’s password database. For more information on the `login` option, refer to “Using the `login` Option With `/etc/ppp/pap-secrets`” on page 549.

1. Assemble a list of all potential trusted callers. Trusted callers are people to be granted permission to call the dial-in server from their remote machines.
2. Verify that each trusted caller already has a UNIX user name and password in the dial-in server's password database.

Note – Verification is particularly important for the sample PAP configuration, which uses the `login` option of PAP to authenticate callers. If you choose not to implement `login` for PAP, the callers' PAP user names do not have to correspond with their UNIX user names. For information on standard `/etc/ppp/pap-secrets`, refer to “`/etc/ppp/pap-secrets` File” on page 546.

Do the following if a potential trusted caller does not have a UNIX user name and password:

- a. Verify with their managers that callers whom you do not know personally have permission to access the dial-in server.
 - b. Create UNIX user names and passwords for these callers in the manner that is directed by your corporate security policy.
3. Become superuser on the dial-in server, and edit the `/etc/ppp/pap-secrets` file. Solaris PPP 4.0 provides a `pap-secrets` file in `/etc/ppp` that contains comments about how to use PAP authentication but no options. You can add the following options at the end of the comments.

```
#
user1      myserver      ""          *
user2      myserver      ""          *
myserver   user2         serverpass  *
```

To use the `login` option of `/etc/ppp/pap-secrets`, you must type the UNIX user name of each trusted caller. Wherever a set of double quotes ("") appears in the third field, the password for the caller is looked up in the server's password database.

The entry `myserver * serverpass *` contains the PAP user name and password for the dial-in server. In Figure 30–3, the trusted caller `user2` requires authentication from remote peers. Therefore, `myserver's` `/etc/ppp/pap-secrets` file contains PAP credentials for use when a link is established with `user2`.

Where to Go From Here

Task	For Instructions
Modify the PPP configuration files to support PAP authentication	"Modifying the PPP Configuration Files for PAP (Dial-in Server)" on page 485
Set up PAP authentication on the dial-out machines of trusted callers	"Configuring PAP Authentication for Trusted Callers (Dial-out Machines)" on page 486

Modifying the PPP Configuration Files for PAP (Dial-in Server)

The tasks in this section explain how to update any existing PPP configuration files to support PAP authentication on the dial-in server.

▼ How to Add PAP Support to the PPP Configuration Files (Dial-in Server)

The procedure uses the PPP configuration files that were introduced in "How to Define Communications Over the Serial Line (Dial-in Server)" on page 471 as examples.

1. Log in to the dial-in server as superuser.

2. Add authentication options to the `/etc/ppp/options` file.

For example, you would add the options in bold to an existing `/etc/ppp/options` file to implement PAP authentication:

```
lock
idle 120
nodefaultroute
name myserver
auth
require-pap
user myserver
remotename user2
login
```

name myserver	Sets myserver as the PAP name of the user on the local machine. If the <code>login</code> option is used, the PAP name must be the same as the user's UNIX user name in the password database.
auth	States that the server must authenticate callers before establishing the link.
require-pap	Requires callers to provide PAP credentials.
user myserver	Defines myserver as the user name of the local machine.
remotename user2	Defines user2 as a peer that requires authentication credentials from the local machine.
login	Specifies that the local machine must use the <code>login</code> option of PAP for authentication wherever <code>login</code> is called for in the <code>/etc/ppp/pap-secrets</code> file.

3. Create an `/etc/ppp/options.ttyname` file, as described in “How to Define Communications Over the Serial Line” on page 463.
4. Set up the `$HOME/.ppprc` file for each remote caller, as explained in “How to Configure Users of the Dial-in Server” on page 470.

Where to Go From Here

Task	For Instructions
Configure PAP authentication credentials for trusted callers of the dial-in server	“Configuring PAP Authentication for Trusted Callers (Dial-out Machines)” on page 486

Configuring PAP Authentication for Trusted Callers (Dial-out Machines)

This section contains tasks for setting up PAP authentication on the dial-out machines of trusted callers. As system administrator, you can set up PAP authentication on the systems before distribution to prospective callers. Or, if the remote callers already have their machines, you can give these callers the tasks in this section.

Configuring PAP for trusted callers involves two tasks:

- Configuring the callers' PAP security credentials

- Configuring the callers' dial-out machines to support PAP authentication

▼ How to Configure PAP Authentication Credentials for the Trusted Callers

This procedure shows how to set up PAP credentials for two trusted callers, one of which requires authentication credentials from remote peers. The steps in the procedure assume that you, the system administrator, are creating the PAP credentials on the trusted callers' dial-out machines.

1. Become superuser on a dial-out machine.

Using the sample PAP configuration that was introduced in Figure 30-3, assume that the dial-out machine belongs to `user1`.

2. Modify the `pap-secrets` database for the caller.

Solaris PPP 4.0 provides an `/etc/ppp/pap-secrets` file that contains helpful comments but no options. You can add the following options to this `/etc/ppp/pap-secrets` file.

```
# user1 myserver pass1 *
```

Note that `user1`'s password `pass1` is passed in readable ASCII form over the link. `myserver` is caller `user1`'s name for the peer.

3. Become superuser on another dial-out machine.

Using the PAP authentication example, assume that this dial-out machine belongs to the caller `user2`.

4. Modify the `pap-secrets` database for the caller.

You can add the next options to the end of the existing `/etc/ppp/pap-secrets` file.

```
# user2 myserver pass2 *
myserver user2 serverpass *
```

In this example, `/etc/ppp/pap-secrets` has two entries. The first entry contains the PAP security credentials that `user2` passes to dial-in server `myserver` for authentication.

`user2` requires PAP credentials from the dial-in server as part of link negotiation. Therefore, the `/etc/ppp/pap-secrets` also contains PAP credentials that are expected from `myserver` on the second line.

Note – Because most ISPs do not supply authentication credentials, the previous scenario might be unrealistic for communications with an ISP.

Where to Go From Here

Task	Instructions
Create PAP credentials for additional callers	"How to Create a PAP Credentials Database (Dial-in Server)" on page 483
Configure a dial-out machine to support PAP authentication	"How to Configure PAP Authentication Credentials for the Trusted Callers" on page 487

Modifying PPP Configuration Files for PAP (Dial-out Machine)

The following tasks explain how to update existing PPP configuration files to support PAP authentication on the dial-out machines of trusted callers.

The procedure uses the following parameters to configure PAP authentication on the dial-out machine that belongs to `user2`, who was introduced in Figure 30–3. `user2` requires incoming callers to authenticate, including calls from `dial-in myserver`.

▼ How to Add PAP Support to the PPP Configuration Files (Dial-out Machine)

This procedure uses the PPP configuration files that were introduced in "How to Define Communications Over the Serial Line" on page 463 as examples. The procedure configures the dial-out machine that belongs to `user2`, as shown in Figure 30–3.

1. **Log in to the dial-out machine as superuser.**
2. **Modify the `/etc/ppp/options` file.**

The next `/etc/ppp/options` file contains options for PAP support, which are shown in bold.

```
#vi /etc/ppp/options
lock
nodefaultroute
name user2
auth
require-pap
```


name user2	Sets <code>user2</code> as the PAP name of the user on the local machine. If the <code>login</code> option is used, the PAP name must be the same as the user's UNIX user name in the password database.
auth	States that the dial-out machine must authenticate callers before establishing the link.
require-pap	Requires peers to provide PAP credentials when returning the call from the dial-out machine.

3. Create an `/etc/ppp/peers/peer-name` file for the remote machine `myserver`.

The next sample shows how to add PAP support to the existing `/etc/ppp/peers/myserver` file that was created in “How to Define the Connection With an Individual Peer” on page 465.

```
# cd /etc/ppp
# mkdir peers
# cd peers
# vi myserver
/dev/cua/a
57600
noipdefault
defaultroute
idle 120
user user2
remotename myserver
connect "chat -U 'mypassword' -f /etc/ppp/mychat"
```

The new options in bold add PAP requirements for peer `myserver`.

user user2	Defines <code>user2</code> as the user name of the local machine
remotename myserver	Defines <code>myserver</code> as a peer that requires authentication credentials from the local machine

Where to Go From Here

Task	For Instructions
Test the PAP authentication setup by calling the dial-in server	Procedures for calling the dial-in server, “How to Call the Dial-in Server” on page 473
Learn more about PAP authentication	“Password Authentication Protocol (PAP)” on page 546

Configuring CHAP Authentication

The tasks in this section explain how to implement authentication on a PPP link by using the Challenge-Handshake Authentication Protocol (CHAP). The tasks use the example that is shown in Figure 30–4 to illustrate a working CHAP scenario for dialing up a private network. Use the instructions as the basis for implementing CHAP authentication at your site.

Before you perform the next procedures, you must have done the following:

- Set up and tested the dial-up link between the dial-in server and dial-out machines that belong to trusted callers
- Obtained superuser permission for the local machine, either dial-in server or dial-out machine

Setting Up CHAP Authentication (Task Maps)

TABLE 33–4 Task Map for CHAP Authentication (Dial-in Server)

Task	Description	For Instructions
1. Assign CHAP secrets to all trusted callers	Create, or have the callers create, their CHAP secrets	"How to Create a CHAP Credentials Database (Dial-in Server)" on page 491
2. Create the chap-secrets database	Add the security credentials for all trusted callers to the <code>/etc/ppp/chap-secrets</code> file	"How to Create a CHAP Credentials Database (Dial-in Server)" on page 491
3. Modify the PPP configuration files	Add options specific to CHAP to the <code>/etc/ppp/options</code> and <code>/etc/ppp/peers/peer-name</code> files	"How to Add CHAP Support to the PPP Configuration Files (Dial-in Server)" on page 492

TABLE 33–5 Task Map for CHAP Authentication (Dial-out Machine)

Task	Description	For Instructions
1. Create the CHAP database for the trusted caller's machine	Create the security credentials for the trusted caller and, if necessary, security credentials for other users who call the dial-out machine, in <code>/etc/ppp/chap-secrets</code> .	"How to Create a CHAP Credentials Database (Dial-in Server)" on page 491
2. Modify the PPP configuration files	Add options specific to CHAP to the <code>/etc/ppp/options</code> file.	"How to Add CHAP Support to the PPP Configuration Files (Dial-out Machine)" on page 494

Configuring CHAP Authentication on the Dial-in Server

The first task in setting up CHAP authentication is modifying the `/etc/ppp/chap-secrets` file. This file contains the CHAP security credentials, including the CHAP secret, that are used to authenticate callers on the link.

Note – UNIX or PAM authentication mechanisms do not work with CHAP. For example, you cannot use the PPP `login` option as described in “How to Create a PAP Credentials Database (Dial-in Server)” on page 483. If your authentication scenario requires PAM or UNIX-style authentication, choose PAP instead.

The next procedure implements CHAP authentication for a dial-in server in a private network. The PPP link is the only connection to the outside world. The only callers who can access the network have been granted permission by managers of the network, possibly including the system administrator.

▼ How to Create a CHAP Credentials Database (Dial-in Server)

1. Assemble a list that contains the user names of all trusted callers. Trusted callers include all people who have been granted permission to call the private network.
2. Assign each user a CHAP secret.

Note – Be sure to choose a good CHAP secret that is not easily guessed. No other restrictions are placed on the CHAP secret’s contents.

The method for assigning CHAP secrets depends on your site’s security policy. Either you have the responsibility for creating the secrets, or the callers must create their own secrets. If you are not responsible for CHAP secret assignment, be sure to get the CHAP secrets that were created by, or for, each trusted caller.

3. Become superuser on the dial-in server, and modify the `/etc/ppp/chap-secrets` file.

Solaris PPP 4.0 includes an `/etc/ppp/chap-secrets` file that contains helpful comments but no options. You can add the following options for the server `CallServe` at the end of the existing `/etc/ppp/chap-secrets` file.

```
account1  CallServe  key123  *
account2  CallServe  key456  *
```

key123 is the CHAP secret for trusted caller account1. key456 is the CHAP secret for trusted caller account2.

Where to Go From Here

Task	For Instructions
Create CHAP credentials for additional trusted callers	"How to Create a CHAP Credentials Database (Dial-in Server)" on page 491
Update the PPP configuration files to support CHAP	"How to Add CHAP Support to the PPP Configuration Files (Dial-in Server)" on page 492
Set up CHAP authentication on the dial-out machines of trusted callers	"Configuring CHAP Authentication for Trusted Callers (Dial-out Machines)" on page 493

Modifying the PPP Configuration Files for CHAP (Dial-in Server)

The task in this section explains how to update existing PPP configuration files to support CHAP authentication on the dial-in server.

▼ How to Add CHAP Support to the PPP Configuration Files (Dial-in Server)

1. **Log in to the dial-in server as superuser.**
2. **Modify the `/etc/ppp/options` file.**
Add the options that are shown in bold for CHAP support.

```
# vi /etc/ppp/options
lock
nodefaultroute
name CallServe
auth
require-chap
```

name CallServe	Defines CallServe as the CHAP name of the user on the local machine, in this instance the dial-in server

auth	Makes the local machine authenticate callers before establishing the link
require-chap	Requires peers to provide CHAP credentials before the link can be established

3. Create the remaining PPP configuration files to support the trusted callers.

See “How to Configure Users of the Dial-in Server” on page 470 and “How to Define Communications Over the Serial Line (Dial-in Server)” on page 471.

Where to Go From Here

Task	For Instructions
Configure CHAP authentication credentials for trusted callers	“How to Create a CHAP Credentials Database (Dial-in Server)” on page 491

Configuring CHAP Authentication for Trusted Callers (Dial-out Machines)

This section contains tasks for setting up CHAP authentication on the dial-out machines of trusted callers. Depending on your site’s security policy, either you or the trusted callers might be responsible for setting up CHAP authentication.

For remote callers to configure CHAP, ensure that the callers’ local CHAP secrets match the callers’ equivalent CHAP secrets in the dial-in server’s `/etc/ppp/chap-secrets` file. Then give the callers the tasks in this section for configuring CHAP.

Configuring CHAP for trusted callers involves two tasks:

- Creating the callers’ CHAP security credentials
- Configuring the callers’ dial-out machines to support CHAP authentication

▼ How to Configure CHAP Authentication Credentials for the Trusted Callers

This procedure shows how to set up CHAP credentials for two trusted callers. The steps in the procedure assume that you, the system administrator, are creating the CHAP credentials on the trusted callers’ dial-out machines.

1. Become superuser on a dial-out machine.

Using the sample CHAP configuration in “Example—Configuration Using CHAP Authentication” on page 452, assume that the dial-out machine belongs to trusted caller `account1`.

2. Modify the chap-secrets database for caller `account1`.

Solaris PPP 4.0 includes an `/etc/ppp/chap-secrets` file that has helpful comments but no options. You can add the following options to the existing `/etc/ppp/chap-secrets` file.

```
# account1 CallServe key123 *
```

`CallServe` is the name for the peer that `account1` is trying to reach. `key123` is the CHAP secret to be used for links between `account1` and `CallServer`.

3. Become superuser on another dial-out machine.

Assume that this machine belongs to caller `account2`.

4. Modify the `/etc/ppp/chap-secrets` database for caller `account2`.

```
# account2 CallServe key456 *
```

Now, `account2` has secret `key456` as its CHAP credentials for use over links to peer `CallServe`.

Where to Go From Here

Task	For Instructions
Create CHAP credentials on the dial-out machines of trusted callers	“How to Create a CHAP Credentials Database (Dial-in Server)” on page 491
Configure a dial-out machine to support CHAP authentication	“How to Configure CHAP Authentication Credentials for the Trusted Callers” on page 493

▼ Adding CHAP to the Configuration Files (Dial-out Machine)

The next task configures the dial-out machine that belongs to caller `account1`, which is introduced in “Example—Configuration Using CHAP Authentication” on page 452.

How to Add CHAP Support to the PPP Configuration Files (Dial-out Machine)

1. Log in to the dial-out machine as superuser.

2. Ensure that the `/etc/ppp/options` file has the following options.

```
# vi /etc/ppp/options
lock
nodefaultroute
```

3. Create an `/etc/ppp/peers/peer-name` file for the remote machine `CallServe`.

```
# mkdir /etc/ppp/peers
# vi CallServe
/dev/cua/a
57600
noipdefault
defaultroute
idle 120
user account1
connect "chat -U 'mypassword' -f /etc/ppp/mychat"
```

The option **user account1** sets `account1` as the CHAP user name to be given to `CallServe`. For a description of the other options in the previous file, see the similar `/etc/ppp/peers/myserver` file in “How to Define the Connection With an Individual Peer” on page 465.

Where to Go From Here

Task	For Instructions
Test CHAP authentication by calling the dial-in server	“How to Call the Dial-in Server” on page 473
Learn more about CHAP authentication	“Challenge-Handshake Authentication Protocol (CHAP)” on page 549

Composed September 5, 2002

Setting Up a PPPoE Tunnel (Tasks)

This chapter contains tasks for setting up the participants on either end of the PPPoE tunnel: the PPPoE client and PPPoE access server. Specific topics include the following:

- “Major Tasks for Setting Up a PPPoE Tunnel (Task Maps)” on page 497
- “Setting Up the PPPoE Client” on page 498
- “Setting Up a PPPoE Access Server” on page 501

The tasks use the scenario that was introduced in “Planning for DSL Support Over a PPPoE Tunnel” on page 454 as an example. For an overview of PPPoE, refer to “Support for DSL Users Through PPPoE” on page 436.

Major Tasks for Setting Up a PPPoE Tunnel (Task Maps)

The following tables list the major tasks for configuring PPPoE clients and the PPPoE access server. To implement PPPoE at your site, you need to set up only your end of the PPPoE tunnel, either the client side or access-server side.

TABLE 34–1 Task Map for Setting Up a PPPoE Client

Task	Description	For Instructions
1. Configure an interface for PPPoE	Define the Ethernet interface to be used for the PPPoE tunnel.	“How to Configure an Interface for a PPPoE Client” on page 499

TABLE 34-1 Task Map for Setting Up a PPPoE Client *(Continued)*

Task	Description	For Instructions
2. Configure information about the PPPoE access server	Define parameters for the access server at the service provider end of the PPPoE tunnel.	"How to Define a PPPoE Access Server Peer" on page 499
3. Set up the PPP configuration files	Define the PPP configuration files for the client, if you have not done so already.	"How to Define Communications Over the Serial Line" on page 463
4. Create the tunnel	Call the access server.	"How to Define a PPPoE Access Server Peer" on page 499

TABLE 34-2 Task Map for Setting Up a PPPoE Access Server

Task	Description	For Instructions
1. Configure an interface for PPPoE	Define the Ethernet interface to be used for the PPPoE tunnel.	"How to Configure the Access Server's Interfaces for PPPoE" on page 501
2. Configure the services that the access server offers	Describe the services that are provided so that these services can be "discovered" by prospective PPPoE clients.	"How to Provide Services to Clients of the Access Server" on page 502
3. Set up the PPP configuration files	Define the PPP configuration files for the client, if you have not done so already.	"Configuring Communications Over the Dial-in Server" on page 471
4. (Optional) Restrict use of an interface	Use PPPoE options and PAP authentication to restrict use of a particular Ethernet interface to certain clients.	"How to Restrict the Use of an Interface to Particular Clients" on page 502

Setting Up the PPPoE Client

To provide PPP to client systems over DSL, you must first configure PPPoE on the interface that is connected to the modem or hub. Then you need to change the PPP configuration files to define the access server on the opposite end of the PPPoE.

Prerequisites for Setting Up the PPPoE Client

Before you set up the PPPoE client, you must have done the following:

- Installed Solaris 8 Update 6 release or subsequent releases on the client machines to use the PPPoE tunnel

- Contacted the service provider for information about its PPPoE access server
- Had the telephone company or service provider assemble the devices that are used by the client machines. These devices include, for example, the DSL modem and the splitter, which the telephone company rather than you might assemble.

▼ How to Configure an Interface for a PPPoE Client

1. Become superuser on the PPPoE client.

2. Add the name of the Ethernet interface with the DSL connection to the `/etc/ppp/pppoe.if` file.

For example, you add the following entry to `/etc/ppp/pppoe.if` for a PPPoE client that uses `hme0` as the network interface that is connected to the DSL modem.

```
hme0
```

For more information about `/etc/ppp/pppoe.if`, go to “`/etc/ppp/pppoe.if` File” on page 555.

3. Configure the interface for PPPoE use.

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

4. (Optional) Verify that the interface is now plumbed for PPPoE.

```
# /usr/sbin/sppptun query
hme0:pppoe
hme0:pppoed
```

You can also use the `/usr/sbin/sppptun` command to manually plumb interfaces for PPPoE. For instructions, refer to “`/usr/sbin/sppptun` Command” on page 556.

▼ How to Define a PPPoE Access Server Peer

You define the access server in the `/etc/ppp/peers/peer-name` file. Many of the options that are used for the access server are also used to define the dial-in server in a dial-up scenario. For a detailed explanation of `/etc/ppp/peers.peer-name`, refer to “`/etc/ppp/peers/peer-name` File” on page 532.

1. Become superuser on the PPPoE client.

2. Define the service provider’s PPPoE access server in the `/etc/ppp/peers/peer-name` file.

For example, the following file, `/etc/ppp/peers/dslserve`, defines the access server `dslserve` at FarISP that is introduced in “Example—Configuration for a PPPoE Tunnel” on page 456.

```
# cat /etc/ppp/peers/dslserve
sppptun
plugin pppoe.so
```

```
connect "/usr/lib/inet/pppoe hme0"
noccip
noauth
user Red
password redsecret
noipdefault
defaultroute
```

For a definition of the options in this file, go to “ /etc/ppp/peers/*peer-name* File for Defining an Access Server Peer” on page 563.

3. Modify the other PPP configuration files on the PPPoE client.

a. Configure /etc/ppp/options as described in the instructions for configuring a dial-out machine in “Configuring the Dial-out Machine” on page 460.

b. Create an /etc/ppp/options.sppptun file. /etc/ppp/options.sppptun defines PPP options for the serial port to which the interface that is plumbed for PPPoE is attached.

You can use any options that are available for the /etc/ppp/options.*ttyname* file that is described in “/etc/ppp/options.*ttyname* Configuration File” on page 529. You must name the file /etc/ppp/options.sppptun because sppptun is the specified device name in the pppd configuration.

4. Ensure that all users can start PPP on the client.

```
# touch /etc/ppp/options
```

5. Test if PPP can run over the DSL line.

```
# pppd debug updetach call dslserve
```

dslserve is the name that is given to the access server at the ISP that is shown in “Example—Configuration for a PPPoE Tunnel” on page 456. The debug updetach option causes debugging information to display in a terminal window.

If PPP is running correctly, the terminal output shows the link becoming active. If PPP still does not run, try the following command to see if the servers are running correctly:

```
# /usr/lib/inet/pppoe -i hme0
```

Where to Go From Here

Task	For Instructions
Configure another PPPoE client	“Setting Up the PPPoE Client” on page 498.
Learn more about PPPoE	“Creating PPPoE Tunnels for DSL Support” on page 554.

Task	For Instructions
Have users of configured PPPoE clients begin running PPP over the DSL line	Instruct the users to type <code>pppd call ISP-server-name</code> and then run an application or service.
Troubleshoot PPPoE and PPP problems	Chapter 35.
Configure a PPPoE access server	“Setting Up a PPPoE Access Server” on page 501.

Setting Up a PPPoE Access Server

If your company is a service provider, you can offer Internet and other services to clients that reach your site through DSL connections. First, you must determine which interfaces on the server to involve in the PPPoE tunnel. Then, you define which services are made available to the users.

▼ How to Configure the Access Server’s Interfaces for PPPoE

1. **Become superuser on the access server.**
2. **Add the name of the Ethernet interfaces that are dedicated to the PPPoE tunnels to the `/etc/ppp/pppoe.if` file.**

For example, you would use the following `/etc/ppp/pppoe.if` file for the access server `dslserve` that is shown in “Example—Configuration for a PPPoE Tunnel” on page 456.

```
# cat /etc/ppp/pppoe.if
hme1
hme2
```

3. **Configure the interfaces for PPPoE use.**

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

4. **(Optional) Verify that interfaces on the server are now plumbed for PPPoE.**

```
# /usr/sbin/sppptun query
hme1:pppoe
hme1:pppoed
hme2:pppoe
hme2:pppoed
```

The previous sample shows that interfaces `hme1` and `hme2` are currently plumbed for PPPoE. You can also use the `/usr/sbin/sppptun` command to manually plumb interfaces for PPPoE. For instructions, refer to “`/usr/sbin/sppptun` Command”

on page 556.

▼ How to Provide Services to Clients of the Access Server

1. **Become superuser on the access server.**
2. **Define global services that are provided by the access server in the `/etc/ppp/pppoe` file.**

The following `/etc/ppp/pppoe` file lists the services that are provided by access server `dslserve`, which was shown in Figure 30–5.

```
device hme1,hme2
service internet
    pppd "proxyarp 192.168.1.1:"
service debugging
    pppd "debug proxyarp 192.168.1.1:"
```

In the file example, Internet service is announced for `dslserve`'s Ethernet interfaces `hme1` and `hme2`. Debugging is turned on for PPP links on the Ethernet interfaces.

3. **Set up the PPP configuration files in the same way that you would for a dial-in server.**

For steps to use, see “Configuring Communications Over the Dial-in Server” on page 471.

4. **Start the `pppoed` daemon.**

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

`pppd` also plumbs the interfaces that are listed in `/etc/ppp/pppoe.if`.

▼ How to Modify an Existing `/etc/ppp/pppoe` File

1. **Become superuser on the access server.**
2. **Modify `/etc/ppp/pppoe`, as needed.**
3. **Cause the `pppoed` daemon to recognize the new services.**

```
# pkill -HUP pppoed
```

▼ How to Restrict the Use of an Interface to Particular Clients

The next procedure shows how to restrict an interface to a group of PPPoE clients. Before performing this task, you need to obtain the real Ethernet MAC addresses of the clients you are assigning to the interface.

Note – Some systems allow you to change the MAC address on the Ethernet interface. You should view this ability as a convenience factor, not a security measure.

Using the example that is shown in “Example—Configuration for a PPPoE Tunnel” on page 456, these steps show how to reserve one of `dslserve`’s interfaces, `hme1`, to clients at MiddleCo.

1. **Configure the access server’s interfaces, as shown in “How to Configure the Access Server’s Interfaces for PPPoE” on page 501.**
2. **Define services, as shown in “How to Provide Services to Clients of the Access Server” on page 502.**
3. **Create entries for clients in the server’s `/etc/ethers` database.**

Here is a sample entry for clients Red, Blue, and Yellow.

```
8:0:20:1:40:30 redether
8:0:20:1:40:10 yellowether
8:0:20:1:40:25 blueether
```

The sample assigns the symbolic names `redether`, `yellowether`, and `blueether` to the Ethernet addresses of clients Red, Yellow, and Blue. The assignment of symbolic names to the MAC addresses is optional.

4. **Restrict services that are provided on a specific interface by defining the following information in the `/etc/ppp/pppoe.device` file.**

In this file, *device* is the name of the device to be defined.

```
# vi /etc/ppp/pppoe.hme1
service internet
    pppd "name dslserve-hme1"
        clients redether,yellowether,blueether
```

`dslserve-hme1` is the access server’s name, which is used in matching entries in the `pap-secrets` file. The `clients` option restricts the use of interface `hme1` to clients with the symbolic Ethernet names `redether`, `yellowether`, and `blueether`.

If you did not define symbolic names for client’s MAC addresses in `/etc/ethers`, you can use the numeric addresses as arguments to the `clients` option. Wildcards are allowed.

For example, you can specify the numeric address `clients 8:0:20:*:*:*`. This address allows access only to clients that are listed in `/etc/ethers` with MAC addresses that begin with the number 8:0:20.

5. **Create the `/etc/ppp/pap-secrets` file for the access server:**

```
# Red          dslserve-hme1  redpasswd    *
Blue          dslserve-hme1  bluepasswd   *
Yellow        dslserve-hme1  yellowpassd  *
```

The entries are the PAP names and passwords of clients that are allowed to run PPP over `dslserve`'s `hme1` interface.

For more information on PAP authentication, see “Configuring PAP Authentication” on page 482.

Where to Go From Here

Task	For Instructions
Learn more about PPPoE	“Creating PPPoE Tunnels for DSL Support” on page 554
Troubleshoot PPPoE and PPP problems	“Diagnosing and Fixing PPPoE Problems” on page 519
Configure a PPPoE client	“Setting Up the PPPoE Client” on page 498
Configure PAP authentication for a client	“Configuring PAP Authentication for Trusted Callers (Dial-out Machines)” on page 486
Configure PAP authentication on a server	“Configuring PAP Authentication on the Dial-in Server” on page 483

Fixing Common Problems (Tasks)

This chapter contains information for troubleshooting common problems that occur with Solaris PPP 4.0. The following topics are covered:

- “Tools for Troubleshooting PPP” on page 506
- “Fixing Network Problems That Affect PPP Performance” on page 509
- “Fixing General Communications Problems” on page 511
- “Fixing PPP Configuration Problems” on page 512
- “Fixing Modem-Related Problems” on page 513
- “Fixing Chat Script-Related Problems” on page 514
- “Fixing Serial-Line Speed Problems” on page 517
- “Fixing Leased-Line Problems” on page 518
- “Diagnosing and Fixing Authentication Problems” on page 519
- “Diagnosing and Fixing PPPoE Problems” on page 519

The sources *PPP Design, Implementation, and Debugging* by James Carlson and the Australian National University’s web site also have detailed advice for PPP troubleshooting. For more information, see “Professional Reference Books About PPP” on page 427 and “Web Sites About PPP” on page 427.

Solving PPP Problems (Task Map)

Use the following task map to quickly access advice and solutions for common PPP problems.

TABLE 35-1 Task Map for Troubleshooting PPP

Task	Definition	For Instructions
Obtain diagnostic information about the PPP link	Use PPP diagnostic tools to get output for troubleshooting	"How to Obtain Diagnostic Information From pppd" on page 507
Obtain debugging information for the PPP link	Use the pppd debug command to generate output for troubleshooting	"How to Turn on PPP Debugging" on page 508
Troubleshoot general problems with the network layer	Identify and fix PPP problems that are network-related by using a series of checks	"How to Diagnose Network Problems" on page 509
Troubleshoot general communications problems	Identify and fix communications problems that affect the PPP link	"How to Diagnose and Fix Communications Problems" on page 511
Troubleshoot configuration problems	Identify and fix problems in the PPP configuration files	"How to Diagnose Problems With the PPP Configuration" on page 513
Troubleshoot modem-related problems	Identify and fix modem problems	"How to Diagnose Modem Problems" on page 514
Troubleshoot chat script-related problems	Identify and fix chat script problems on a dial-out machine	"How to Obtain Debugging Information for Chat Scripts" on page 515
Troubleshoot serial-line speed problems	Identify and fix line speed problems on a dial-in server	"How to Diagnose and Fix Serial-Line Speed Problems" on page 517
Troubleshoot common problems for leased lines	Identify and fix performance problems on a leased line	"Fixing Leased-Line Problems" on page 518
Troubleshoot problems related to authentication	Identify and fix problems related to the authentication databases	"Diagnosing and Fixing Authentication Problems" on page 519
Troubleshoot problem areas for PPPoE	Use PPP diagnostic tools to obtain output for identifying and fixing PPPoE problems	"How to Obtain Diagnostic Information for PPPoE" on page 519

Tools for Troubleshooting PPP

PPP links generally have three major areas of failure:

- Failure of the link to be established
- Poor performance of the link during regular usage
- Problems that can be traced to the networks on either side of the link

The easiest way to find out if PPP works is to run a command over the link. Run a command such as `ping` or `traceroute` to a host on the peer's network. Then observe the results. However, you should use PPP and UNIX debugging tools to monitor performance of an established link or to troubleshoot a problematic link.

This section explains how to obtain diagnostic information from `pppd` and its associated log files. The remaining sections in this chapter describe common problems with PPP that you can discover and fix with the aid of the PPP troubleshooting tools.

▼ How to Obtain Diagnostic Information From `pppd`

The next procedure shows how to view the current operation of a link on the local machine.

1. Become superuser on the local machine.

2. Run `pppd` with the serial device configured for PPP as the argument:

```
# pppd /dev/ttyname debug updetach
```

The next examples show the resulting displays for a dial-up link and a leased-line link when `pppd` runs in the foreground. If you run `pppd debug` in the background, the output that is produced is sent to the `/etc/ppp/connect-errors` file.

EXAMPLE 35-1 Output From a Properly Operating Dial-up Link

```
# pppd /dev/cua/b debug updetach
have route to 0.0.0.0/0.0.0.0 via 172.21.0.4
serial speed set to 230400 bps
Using interface sppp0
Connect: sppp0 <--> /dev/cua/b
sent [LCP ConfReq id=0x7b <asynctest 0x0> <magic 0x73e981c8> <pcomp> <accomp>]
rcvd [LCP Ident id=0x79 magic=0x0 "ppp-2.4.0b1 (Sun Microsystems, Inc., Dec 6
2000 09:36:22)"]
Peer Identification: ppp-2.4.0b1 (Sun Microsystems, Inc., Dec 6 2000 09:36:22)
rcvd [LCP ConfReq id=0x7b <asynctest 0x0>]
sent [LCP Ident id=0x7c magic=0x0 "ppp-2.4.0b1 (Sun Microsystems, Inc., Nov 15
2000 09:38:33)"]
sent [LCP ConfReq id=0x7d <magic 0x73e981c8> <pcomp> <accomp>]
rcvd [LCP ConfAck id=0x7d <magic 0x73e981c8> <pcomp> <accomp>]
rcvd [LCP ConfAck id=0x78 <magic 0xdd4ad820> <pcomp> <accomp>]
sent [LCP ConfAck id=0x78 <magic 0xdd4ad820> <pcomp> <accomp>]
sent [LCP Ident id=0x7e magic=0x73e981c8 "ppp-2.4.0b1 (Sun Microsystems, Inc.,
Nov 15 2000 09:38:33)"]
sent [IPCP ConfReq id=0x3d <addr 0.0.0.0> <compress VJ 0f 01>]
rcvd [LCP Ident id=0x7a magic=0xdd4ad820 "ppp-2.4.0b1 (Sun Microsystems, Inc.,
Dec 6 2000 09:36:22)"]
Peer Identification: ppp-2.4.0b1 (Sun Microsystems, Inc., Dec 6 2000 09:36:22)
rcvd [IPCP ConfReq id=0x92 <addr 10.0.0.1> <compress VJ 0f 01>]
sent [IPCP ConfAck id=0x92 <addr 10.0.0.1> <compress VJ 0f 01>]
rcvd [IPCP ConfNak id=0x3d <addr 10.0.0.2>]
sent [IPCP ConfReq id=0x3e <addr 10.0.0.2> <compress VJ 0f 01>]
```

Composed September 5, 2002

EXAMPLE 35-1 Output From a Properly Operating Dial-up Link (Continued)

```
rcvd [IPCP ConfAck id=0x3e <addr 10.0.0.2> <compress VJ Of 01>]
local IP address 10.0.0.2
remote IP address 10.0.0.1
```

EXAMPLE 35-2 Output From a Properly Operating Leased-Line Link

```
# pppd /dev/se_hdlc1 default-asynctmap debug updetach
pppd 2.4.0b1 (Sun Microsystems, Inc., Oct 24 2001 07:13:18) started by root, uid 0
synchronous speed appears to be 0 bps
init option: '/etc/ppp/peers/syncinit.sh' started (pid 105122)
Serial port initialized.
synchronous speed appears to be 64000 bps
Using interface sppp0
Connect: sppp0 <--> /dev/se_hdlc1
sent [LCP ConfReq id=0xe9 <magic 0x474283c6><pcomp> <accomp>]
rcvd [LCP ConfAck id=0xe9 <magic 0x474283c6><pcomp> <accomp>]
rcvd [LCP ConfReq id=0x22 <magic 0x8e3a53ff><pcomp> <accomp>]
sent [LCP ConfReq id=0x22 <magic 0x8e3a53ff><pcomp> <accomp>]
sent [LCP Ident id=0xea magic=0x474283c6 "ppp-2.4.0b1 (Sun Microsystems, Inc., Oct
22 2001 14:31:44)"]
sent [IPCP ConfReq id=0xf7 <addr 0.0.0.0> <compress VJ Of 01>]]
sent [CCP ConfReq id=0x3f <deflate 15> <deflate(old#) 15> <bsd v1 15>]
rcvd [LCP Ident id=0x23 magic=0x8e3a53ff "ppp-2.4.0b1 (Sun Microsystems, Inc., Oct
22 2001 14:31:44)"]
Peer Identification: ppp-2.4.0b1 (Sun Microsystems, Inc., Oct 22 2001 14:31:44)
rcvd [IPCP ConfReq id=0x25 <addr 10.0.0.1> <compress VJ Of 01>]
sent [IPCP ConfAck id=0x25 <addr 10.0.0.1> <compress VJ Of 01>]
rcvd [CCP ConfReq id=0x3 <deflate 15> <deflate(old#) 15> <bsd v1 15>]
sent [CCP ConfAck id=0x3 <deflate 15> <deflate(old#) 15> <bsd v1 15>]
rcvd [IPCP ConfNak id=0xf8 <addr 10.0.0.2>]
rcvd [IPCP ConfReq id=0xf7 <addr 10.0.0.2> <compress VJ Of 01>]
rcvd [CCP ConfAck id=0x3f <deflate 15> <deflate(old#) 15> <bsd v1 15>]
Deflate (15) compression enabled
rcvd [IPCP ConfAck id=0xf8 <addr 10.0.0.2> <compress VJ Of 01>]
local IP address 10.0.0.2
remote IP address 10.0.0.1
```

▼ How to Turn on PPP Debugging

The next task shows how to use the `pppd` command to obtain debugging information.

Note – You only need to perform step 1 through step 3 once for each host. Thereafter, you can proceed to step 4 to turn on debugging for the host.

1. Create a log file to hold output from `pppd`.

```
% touch /var/log/pppdebug
```

2. Add the following syslog facilities for pppd in `/etc/syslog.conf`.
`daemon.debug;local2.debug           /var/log/pppdebug`
 3. Restart syslogd.
`% pkill -HUP -x syslogd`
 4. Turn on debugging for calls to a particular peer by using the following syntax of pppd.
`% pppd debug call peer-name`
peer-name must be the name of a file in the `/etc/ppp/peers` directory.
 5. View the contents of the log file.
`% tail -f /var/log/pppdebug`
For an example of a log file, see Example 35–3.
-

Fixing Network Problems That Affect PPP Performance

If the PPP link becomes active but few hosts on the remote network are reachable, a network problem could be indicated. This section explains how to isolate and fix network problems that affect a PPP link.

▼ How to Diagnose Network Problems

1. Become superuser on the local machine. Then shut down the problematic link.
2. Disable any optional protocols in the configuration files by adding the following options to your PPP configuration:
`noccp novj nopcomp noaccomp default-asynmap`
These options provide the simplest uncompressed PPP that is available. Try to invoke these options as arguments to pppd on the command line. If you can reach the previously unreachable hosts, add the options in either of the following places.
 - `/etc/ppp/peers/peer-name`, after the `call` option
 - `/etc/ppp/options`, ensuring that the options apply globally
3. Call the remote peer. Then enable debugging features.
`% pppd debug call peer-name`
4. Obtain verbose logs from the chat program by using the `-v` option of chat.

For example, use the following format in any PPP configuration file:

```
connect 'chat -v -f /etc/ppp/chatfile'
```

/etc/ppp/chatfile represents the name of your chat file.

5. Try to re-create the problem by using Telnet or other applications to reach the remote hosts.

Observe the debugging logs. If you still cannot reach remote hosts, the PPP problem might be network-related.

6. Verify that the IP addresses of the remote hosts are registered Internet addresses.

Some organizations assign internal IP addresses that are known within the local network but cannot be routed to the Internet. If the remote hosts are within your company, you must set up a name-to-address translation (NAT) server or proxy server to reach the Internet. If the remote hosts are not within your company, you should report the problem to the remote organization.

7. Examine the routing tables.

a. Check the routing tables on both the local machine and the peer.

b. Check the routing tables for any routers that are in the path from the peer to the remote system. Also check the routing tables for any routers on the path back to the peer.

Ensure that the intermediate routers have not been misconfigured. Often the problem can be found in the path back to the peer.

8. (Optional) If the machine is a router, check the optional features.

```
# ndd -set /dev/ip ip_forwarding 1
```

For more information about ndd, refer to the ndd(1M) man page.

9. Check the statistics that are obtained from netstat -s and similar tools.

For complete details on netstat, refer to the netstat(1M) man page.

a. Run statistics on the local machine.

b. Call the peer.

c. Observe the new statistics that are generated by netstat -s.

You can use the messages that are generated by netstat -s to fix the network problems that are shown in the next table.

TABLE 35-2 Common Network Problems That Affect PPP

Message	Problem	Solution
IP packets not forwardable	The local host is missing a route.	Add the missing route to the local host's routing tables.

TABLE 35-2 Common Network Problems That Affect PPP (Continued)

Message	Problem	Solution
ICMP input destination unreachable	The local host is missing a route.	Add the missing route to the local host's routing tables.
ICMP time exceeded	Two routers are forwarding the same destination address to each other, causing the packet to bounce back and forth until the time-to-live (TTL) value is exceeded.	Use <code>tracert</code> to find the source of the routing loop, and then contact the administrator of the router in error. For information about <code>tracert</code> , refer to the <code>tracert(1M)</code> man page.
IP packets not forwardable	The local host is missing a route.	Add the missing route to the local host's routing table.
ICMP input destination unreachable	The local host is missing a route.	Add the missing route to the local host's routing tables.

10. Check the DNS configuration.

A faulty name service configuration causes applications to fail because IP addresses cannot be resolved.

You can find information for fixing name service problems in “DNS Troubleshooting (Reference)” in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

Fixing General Communications Problems

Communications problems occur when the two peers cannot successfully establish a link. Sometimes these problems are actually negotiation problems that are caused by incorrectly configured chat scripts. This section contains information for clearing communications problems. For clearing negotiation problems that are caused by a faulty chat script, see Table 35-5.

▼ How to Diagnose and Fix Communications Problems

1. Become superuser on the local machine. Then call the peer.
2. Call the remote peer. Then enable debugging features.

```
% pppd debug call peer-name
```

You might need to obtain debugging information from the peer in order to fix certain communications problems.

3. Check the resulting logs for the communications problems in the next table.

TABLE 35-3 General Communications Problems That Affect PPP

Symptom	Problem	Solution
too many Configure-Requests	One peer cannot hear the other peer.	Check for the following problems: ■ The machine or modem might have faulty cabling. ■ The modem configuration might have incorrect bit settings. Or, the configuration might have broken flow control. ■ The chat script might have failed. In this situation, see Table 35-5.
The <code>pppd</code> debug output shows that LCP starts, but higher-level protocols fail or show CRC errors.	The asynchronous control character map (ACCM) is incorrectly set.	Use the <code>default-async</code> option to set the ACCM to the standard default of FFFFFFFF. First try to use <code>default-async</code> as an option to <code>pppd</code> on the command line. If the problem clears, then add <code>default-async</code> to <code>/etc/ppp/options</code> or to <code>/etc/ppp/peers/peer-name</code> after the call option.
The <code>pppd</code> debug output shows that IPCP comes up but terminates immediately.	IP addresses might be incorrectly configured.	- Check the chat script to verify whether the script has incorrect IP addresses. - If the chat script is okay, request debug logs for the peer, and check IP addresses in the peer logs.
The link exhibits very poor performance.	The modem might be incorrectly configured, with flow-control configuration errors, modem setup errors, and incorrectly configured DTE rates.	Check the modem configuration. Adjust the configuration, if necessary.

Fixing PPP Configuration Problems

Some PPP problems can be traced to problems in the PPP configuration files. This section contains information for isolating and fixing general configuration problems.

▼ How to Diagnose Problems With the PPP Configuration

1. Become superuser on the local machine.
2. Call the remote peer. Then enable debugging features.

```
% pppd debug call peer-name
```
3. Check the resulting log for the configuration problems that are listed in the next table.

TABLE 35-4 Common PPP Configuration Problems

Symptom	Problem	Solution
pppd debug output contains the error message "Could not determine remote IP address."	The <code>/etc/ppp/peers/peer-name</code> file does not have an IP address for the peer. The peer does not provide an IP address during link negotiation.	Supply an IP address for the peer on the pppd command line or in <code>/etc/ppp/peers/peer-name</code> by using the following format: :10.0.0.10
pppd debug output shows that CCP data compression has failed. The output also indicates that the link is dropped.	The peers' PPP compression configurations might be in conflict.	Disable CCP compression by adding the <code>noccp</code> option to <code>/etc/ppp/options</code> on one of the peers.

Fixing Modem-Related Problems

Modems can be major problem areas for a dial-up link. The most common indicator of problems with the modem configuration is no response from the peer. However, you might have difficulties when determining if a link problem is indeed the result of modem configuration problems.

For basic modem troubleshooting suggestions, refer to "Troubleshooting Terminal and Modem Problems" in *System Administration Guide: Advanced Administration*. Modem manufacturers' documentation and web sites also contain solutions for problems with their particular equipment. This section contains more suggestions for isolating and fixing modem problems.

▼ How to Diagnose Modem Problems

The following steps help determine whether a faulty modem configuration causes link problems.

1. **Call the peer with debugging turned on, as explained in “How to Turn on PPP Debugging” on page 508.**

2. **Display the resulting `/var/log/pppdebug` log.**

Either of the following symptoms in the output can indicate a faulty modem configuration:

- No “recvd” messages have come from the peer.
- The output contains LCP messages from the peer, but the link fails with too many LCP Configure Requests messages that are sent by the local machine.

These messages indicate that the local machine can hear the peer, but the peer cannot hear the local machine.

- The link terminates with a SIGHUP signal.

3. **Use ping to send packets of various sizes over the link.**

For complete details about ping, refer to the ping(1M) man page.

If small packets are received but larger packets are dropped, modem problems are indicated.

4. **Check for errors on interface `sppp0`:**

```
% netstat -ni
Name  Mtu  Net/Dest  Address      IpKts    Ierrs  OpKts    Oerrs  Collis  Queue
lo0    8232  127.0.0.0  127.0.0.1    826808   0      826808   0      0      0
hme0   1500  172.21.0.0 172.21.3.228 13800032 0      1648464  0      0      0
sppp0  1500  10.0.0.2   10.0.0.1     210      0      128     0      0      0
```

If interface errors increase over time, the modem configuration might have problems.

Fixing Chat Script-Related Problems

Chat scripts are trouble-prone areas for dial-up links. This section contains a procedure for obtaining debugging information from chat and suggestions for clearing common problems.

▼ How to Obtain Debugging Information for Chat Scripts

1. Become superuser on the dial-out machine.
2. Edit the `/etc/ppp/peers/peer-name` file for the peer to be called.
3. Add `-v` as an argument to the `chat` command that is specified in `connect` option.

```
connect "/usr/bin/chat -v -f /etc/ppp/chat-script-name"
```

4. View chat script errors in the file `/etc/ppp/connect-errors`.

The following is the main error that is seen with `chat`.

```
Oct 31 08:57:13 deino chat[107294]: [ID 702911 local2.info] expect (CONNECT)
Oct 31 08:57:58 deino chat[107294]: [ID 702911 local2.info] alarm
Oct 31 08:57:58 deino chat[107294]: [ID 702911 local2.info] Failed
```

The example shows timeout while waiting for a `(CONNECT)` string. When `chat` fails, you get the following message from `pppd`:

```
Connect script failed
```

The next table lists common chat script errors and suggestions for fixing the errors.

TABLE 35-5 Common Chat Script Problems

Symptom	Problem	Solution
pppd debug output contains Connect script failed	Your chat script supplies a user name and ssword. ogin: <i>user-name</i> ssword: <i>password</i> However, the peer that you intended to connect to does not prompt for this information.	1. Delete the login and password from the chat script. 2. Try to call the peer again. 3. If you still get the message, call the ISP. Ask the ISP for the correct login sequence.
The <code>/usr/bin/chat -v</code> log contains the message: "expect (login:) " alarm read timed out	Your chat script supplies a user name and ssword. ogin: <i>pppuser</i> ssword: <i>\q\U</i> However, the peer that you intend to connect to does not prompt for this information.	1. Delete the login and password from the chat script. 2. Try to call the peer again. 3. If you still get the message, call the ISP. Ask the ISP for the correct login sequence.
pppd debug output contains possibly looped-back	The local machine or its peer is hanging at the command line and not running PPP. An incorrectly configured login name and password are in the chat script.	1. Delete the login and password from the chat script. 2. Try to call the peer again. 3. If you still get the message, call the ISP. Ask for the correct login sequence.

TABLE 35-5 Common Chat Script Problems (Continued)

Symptom	Problem	Solution
pppd debug output shows that LCP activates, but the link terminates soon afterward.	The password in the chat script might be incorrect.	1. Ensure that you have the correct password for the local machine. 2. Check the password in the chat script. Fix the password if incorrect. 3. Try to call the peer again. 4. If you still get the message, call the ISP. Ask the ISP for the correct login sequence.
Text from the peer begins with a tilde (~).	Your chat script supplies a user name and ssword. <pre>ogin: pppuser ssword: \q\U</pre> However, the peer that you intend to connect to does not prompt for this information.	1. Delete the login and password from the chat script. 2. Try to call the peer again. 3. If you still get the message, call the ISP. Request the correct login sequence.
The modem hangs.	Your chat script contains the following line to force the local machine to wait for the CONNECT message from the peer: <pre>CONNECT "</pre>	Use the following line when you want the chat script to wait for CONNECT from the peer: <pre>CONNECT \c</pre> End the chat script with ~ \c.
pppd debug output contains LCP: timeout sending Config-Requests	Your chat script contains the following line to force the local machine to wait for the CONNECT message from the peer: <pre>CONNECT "</pre>	Use the following line when you want the chat script to wait for CONNECT from the peer: <pre>CONNECT \c</pre> End the chat script with ~ \c .
pppd debug output contains: Serial link is not 8-bit clean	Your chat script contains the following line to force the local machine to wait for the CONNECT message from the peer: <pre>CONNECT "</pre>	Use the following line when you want the chat script to wait for CONNECT from the peer: <pre>CONNECT \c</pre> End the chat script with ~ \c.
pppd debug output contains: Loopback detected	Your chat script contains the following line to force the local machine to wait for the CONNECT message from the peer: <pre>CONNECT "</pre>	Use the following line when you want the chat script to wait for CONNECT from the peer: <pre>CONNECT \c</pre> End the chat script with ~ \c.

TABLE 35-5 Common Chat Script Problems (Continued)

Symptom	Problem	Solution
pppd debug output contains SIGHUP	Your chat script contains the following line to force the local machine to wait for the CONNECT message from the peer: CONNECT "	Use the following line when you want the chat script to wait for CONNECT from the peer: CONNECT \c End the chat script with ~ \c.

Fixing Serial-Line Speed Problems

Dial-in servers can experience problems because of conflicting speed settings. The next procedure helps you to isolate the cause of the link problem to conflicting serial-line speeds.

The following behaviors cause speed problems:

- You invoked PPP through a program such as `/bin/login` and specified the speed of the line.
- You started PPP from `mgetty` and accidentally supplied the bit rate.

`pppd` changes the speed that was originally set for the line to the speed that was set by `/bin/login` or `mgetty`. As a result, the line fails.

▼ How to Diagnose and Fix Serial-Line Speed Problems

1. **Log in to the dial-in server. Call the peer with debugging enabled.**

If you need instructions, see “How to Turn on PPP Debugging” on page 508.

2. **Display the resulting `/var/log/pppdebug` log.**

Check the output for the following message:

```
LCP too many configure requests
```

This message indicates that the speeds of serial lines that were configured for PPP might potentially be in conflict.

3. **Check if PPP is invoked through a program such as `/bin/login` and the line speed that was set.**

In such a situation, `pppd` changes the originally configured line speed to the speed that is specified in `/bin/login`.

4. Check if a user started PPP from the `mgetty` command and accidentally specified a bit rate.

This action also causes serial-line speeds to conflict.

5. Fix the conflicting serial-line speed problem as follows:

- a. Lock the DTE rate on the modem.
- b. Do not use autobaud.
- c. Do not change the line speed after configuration.

Fixing Leased-Line Problems

The most common problem with leased lines is poor performance. In most situations, you need to work with the telephone company to fix the problem.

TABLE 35-6 Common Leased-Line Problems

Symptom	Problem	Solution
The link does not start.	CSU bipolar violations (CSU BPVs) can be the cause. One end of the link is set up for AMI lines. The other end is set up for ESF bit 8 zero substitute (B8Zs).	If you are in the United States or Canada, you can directly fix this problem from the menu of the CSU/DSU. Check the CSU/DSU manufacturer's documentation for details. In other locales, the provider might be responsible for fixing CSU BPVs.
The link has poor performance.	The <code>pppd</code> debug output shows CRC errors when sustained traffic is on the link. Your line might have a clocking problem, caused by misconfigurations between the telephone company and your network.	Contact the telephone company to ensure that "loop clocking" is in use. On some unstructured leased lines, you might have to supply clocking. North American users should use loop clocking.

Diagnosing and Fixing Authentication Problems

TABLE 35-7 General Authentication Problems

Symptom	Problem	Solution
pppd debug output shows the message <code>Peer is not authorized to use remote address address</code> .	You are using PAP authentication, and the IP address for the remote peer is not in the <code>/etc/ppp/pap-secrets</code> file.	Add an asterisk (*) after the entry for the peer in the <code>/etc/ppp/pap-secrets</code> file.
pppd debug output shows that LCP starts but terminates shortly afterward.	The password might be incorrect in the database for the particular security protocol.	Check the password for the peer in the <code>/etc/ppp/pap-secrets</code> or <code>/etc/ppp/chap-secrets</code> file.

Diagnosing and Fixing PPPoE Problems

You can use PPP and standard UNIX utilities to identify problems with PPPoE. This section explains how to obtain debugging information for PPPoE. The section also contains suggestions for fixing PPPoE-related problems.

How to Obtain Diagnostic Information for PPPoE

When you suspect that PPPoE is the cause of problems on a link, use the following diagnostic tools to obtain troubleshooting information.

1. **Become superuser on the machine that runs the PPPoE tunnel, either PPPoE client or PPPoE access server.**
2. **Turn on debugging, as explained in the procedure “How to Turn on PPP Debugging” on page 508.**
3. **View the contents of the log file `/var/log/pppdebug`.**

The following example shows part of a log file that was generated for a link with a PPPoE tunnel.

EXAMPLE 35-3 Log File for a Link With a PPPoE Tunnel

```
Sep  6 16:28:45 enyo pppd[100563]: [ID 702911 daemon.info] Plugin
pppoe.so loaded.
```

Composed September 5, 2002

EXAMPLE 35-3 Log File for a Link With a PPPoE Tunnel (Continued)

```
Sep  6 16:28:45 enyo pppd[100563]: [ID 860527 daemon.notice] pppd
2.4.0b1 (Sun Microsystems, Inc.,
Sep  5 2001 10:42:05) started by troot, uid 0
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.debug] connect option:
'/usr/lib/inet/pppoe
-v hme0' started (pid 100564)
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.info] Serial connection established.
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.info] Using interface sPPP0
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.notice] Connect: sPPP0
<--> /dev/sppptun
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.debug] /etc/ppp/pap-secrets
is apparently empty
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.debug] /etc/ppp/chap-secrets
is apparently empty
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.debug] sent
[LCP ConfReq id=0xef <mru 1492>
asynctest 0x0 <magic 0x77d3e953><pcomp><acomp>
Sep  6 16:28:46 enyo pppd[100563]: [ID 702911 daemon.debug] rcvd
[LCP ConfReq id=0x2a <mru 1402>
asynctest 0x0 <magic 0x9985f048><pcomp><acomp>
```

If the debugging output does not help you isolate the problem, continue with this procedure.

4. Get diagnostic messages from PPPoE.

```
# pppd connect "/usr/lib/inet/pppoe -v interface-name"
```

pppoe sends diagnostic information to the `stderr`. If you run `pppd` in the foreground, the output appears on the screen. If `pppd` runs in the background, the output is sent to `/etc/ppp/connect-errors`.

The next example shows the messages that are generated as the PPPoE tunnel is negotiated.

EXAMPLE 35-4 PPPoE Diagnostic Messages

```
Connect option: '/usr/lib/inet/pppoe -v hme0' started (pid 100564)
/usr/lib/inet/pppoe: PPPoE Event Open (1) in state Dead (0): action SendPADI (2)
/usr/lib/inet/pppoe: Sending PADI to ff:ff:ff:ff:ff:ff: 18 bytes
/usr/lib/inet/pppoe: PPPoE State change Dead (0) -> InitSent (1)
/usr/lib/inet/pppoe: Received Active Discovery Offer from 8:0:20:cd:c1:2/hme0:pppoe
/usr/lib/inet/pppoe: PPPoE Event rPADO+ (5) in state InitSent (1): action SendPADR+ (5)
/usr/lib/inet/pppoe: Sending PADR to 8:0:20:cd:c1:2: 22 bytes
/usr/lib/inet/pppoe: PPPoE State change InitSent (1) -> ReqSent (3)
/usr/lib/inet/pppoe: Received Active Discovery Session-confirmation from
8:0:20:cd:c1:2/hme0:pppoe
/usr/lib/inet/pppoe: PPPoE Event rPADS (7) in state ReqSent (3): action Open (7)
/usr/lib/inet/pppoe: Connection open; session 0002 on hme0:pppoe
/usr/lib/inet/pppoe: PPPoE State change ReqSent (3) -> Convers (4)
/usr/lib/inet/pppoe: connected
```


Composed September 5, 2002

If the diagnostic messages do not help you isolate the problem, continue with this procedure.

5. Run snoop. Then save the trace to a file.

For information about snoop, refer to the snoop(1M) man page.

```
# snoop -o pppoe-trace-file
```

6. View the snoop trace file.

```
# snoop -i pppoe-trace-file -v pppoe
```

EXAMPLE 35-5 PPPoE snoop trace

```
ETHER: ----- Ether Header -----
ETHER:
ETHER: Packet 1 arrived at 6:35:2.77
ETHER: Packet size = 32 bytes
ETHER: Destination = ff:ff:ff:ff:ff:ff, (broadcast)
ETHER: Source      = 8:0:20:78:f3:7c, Sun
ETHER: Ethertype = 8863 (PPPoE Discovery)
ETHER:
PPPoE: ----- PPP Over Ethernet -----
PPPoE:
PPPoE: Version = 1
PPPoE: Type = 1
PPPoE: Code = 9 (Active Discovery Initiation)
PPPoE: Session Id = 0
PPPoE: Length = 12 bytes
PPPoE:
PPPoE: ----- Service-Name -----
PPPoE: Tag Type = 257
PPPoE: Tag Length = 0 bytes
PPPoE:
PPPoE: ----- Host-Uniq -----
PPPoE: Tag Type = 259
PPPoE: Tag Length = 4 bytes
PPPoE: Data = 0x00000002
PPPoE:
.
.
.
ETHER: ----- Ether Header -----
ETHER:
ETHER: Packet 5 arrived at 6:35:2.87
ETHER: Packet size = 60 bytes
ETHER: Destination = 8:0:20:78:f3:7c, Sun)
ETHER: Source      = 0:2:fd:39:7f:7,
ETHER: Ethertype = 8864 (PPPoE Session)
ETHER:
PPPoE: ----- PPP Over Ethernet -----
PPPoE:
PPPoE: Version = 1
PPPoE: Type = 1
PPPoE: Code = 0 (PPPoE Session)
```

Composed September 5, 2002

EXAMPLE 35-5 PPPoE snoop trace *(Continued)*

```
PPPoE: Session Id = 24383
PPPoE: Length = 20 bytes
PPPoE:
PPP: ----- Point-to-Point Protocol -----
PPP:
PPP-LCP: ----- Link Control Protocol -----
PPP-LCP:
PPP-LCP: Code = 1 (Configure Request)
PPP-LCP: Identifier = 80
PPP-LCP: Length = 18
```

Solaris PPP 4.0 Reference

This chapter provides detailed conceptual information about Solaris PPP 4.0. Topics include the following:

- “Using PPP Options in Files and on the Command Line” on page 523
- “Configuring User-Specific Options” on page 531
- “Specifying Information for Communicating With the Dial-in Server” on page 532
- “Configuring Modems for a Dial-up Link” on page 535
- “Defining the Conversation on the Dial-up Link” on page 536
- “Authenticating Callers on a Link” on page 545
- “Creating an IP Addressing Scheme for Callers” on page 552
- “Creating PPPoE Tunnels for DSL Support” on page 554

Using PPP Options in Files and on the Command Line

Solaris PPP 4.0 contains a large set of options, which you use to define your PPP configuration. You use these options in the PPP configuration files, or on the command line, or by using a combination of files and command-line options. This section contains detailed information about the use of PPP options in configuration files and as arguments to PPP commands.

Where to Define PPP Options

Solaris PPP 4.0 configuration is very flexible. You can define PPP options in the following places:

- PPP configuration files

- PPP commands that are issued on the command line
- A combination of both places

The next table lists the PPP configuration files and commands.

TABLE 36-1 Summary of PPP Configuration Files and Commands

File or Command	Definition	For Information
<code>/etc/ppp/options</code>	A file that contains characteristics that apply by default to all PPP links on the system, for example, whether the machine requires peers to authenticate themselves. If this file is absent, non-root users are prohibited from using PPP.	<code>/etc/ppp/options</code> Configuration File" on page 528
<code>/etc/ppp/options.ttyname</code>	A file that describes the characteristics of all communications over the serial port <i>ttyname</i> .	<code>/etc/ppp/options.ttyname</code> Configuration File" on page 529
<code>/etc/ppp/peers</code>	Directory that usually contains information about peers with which a dial-out machine connects. Files in this directory are used with the <code>call</code> option of the <code>pppd</code> command.	"Specifying Information for Communicating With the Dial-in Server" on page 532
<code>/etc/ppp/peers/peer-name</code>	A file that contains characteristics of the remote peer <i>peer-name</i> . Typical characteristics include the remote peer's phone number and chat script for negotiating the link with the peer.	<code>/etc/ppp/peers/peer-name</code> File" on page 532
<code>/etc/ppp/pap-secrets</code>	A file that contains the necessary security credentials for Password Authentication Protocol (PAP) authentication.	<code>/etc/ppp/pap-secrets</code> File" on page 546
<code>/etc/ppp/chap-secrets</code>	A file that contains the necessary security credentials for Challenge-Handshake Authentication Protocol (CHAP) authentication.	<code>/etc/ppp/chap-secrets</code> File" on page 550
<code>~/.ppprc</code>	File in the home directory of a PPP user, most often used with dial-in servers. This file contains specific information about each user's configuration.	"Configuring User-Specific Options" on page 531
<code>pppd options</code>	Command and options for initiating a PPP link and describing its characteristics.	"How PPP Options Are Processed" on page 525

Refer to the `pppd(1M)` man page for details on the PPP files. `pppd(1M)` also includes comprehensive descriptions of all options that are available to the `pppd` command. Sample templates for all the PPP configuration files are available in `/etc/ppp`.

How PPP Options Are Processed

All Solaris PPP 4.0 operations are handled by the `pppd` daemon, which starts when a user runs the `pppd` command. When a user calls a remote peer, the following occurs:

1. The `pppd` daemon parses the following:
 - `/etc/ppp/options`
 - `$HOME/.ppprc`
 - Any files that are opened by the `file` or `call` option in `/etc/ppp/options` and `$HOME/.ppprc`
2. `pppd` scans the command line to determine the device in use. The daemon does not yet interpret any options that are encountered.
3. `pppd` tries to discover the serial device to use by using the following criteria:
 - If a serial device is specified on the command line, or a previously processed configuration file, `pppd` uses the name of that device.
 - If no serial device is named, then `pppd` searches for the `notty`, `pty`, or `socket` option on the command line. If one of these options is specified, `pppd` assumes that no device name exists.
 - Otherwise, if `pppd` discovers that standard input is attached to a `tty`, then the name of the `tty` is used.
 - If `pppd` still cannot find a serial device, `pppd` terminates the connection and issues an error.
4. `pppd` then checks for the existence of the `/etc/ppp/options.ttyname` file. If the file is found, `pppd` parses the file.
5. `pppd` processes any options on the command line.
6. `pppd` negotiates the Link Control Protocol (LCP) to set up the link.
7. (Optional) If authentication is required, `pppd` reads `/etc/ppp/pap-secrets` or `/etc/ppp/chap-secrets` to authenticate the opposite peer.

The file `/etc/ppp/peers/peer-name` is read when the `pppd` daemon encounters the option `call peer-name` on the command line or in the other configuration files.

How PPP Configuration File Privileges Work

Solaris PPP 4.0 configuration includes the concept of *privileges*. Privileges determine the precedence of configuration options, particularly when the same option is invoked in more than one place. An option that is invoked from a privileged source takes precedence over the same option that is invoked from a non-privileged source.

User Privileges

The only privileged user is superuser (`root`), with the UID of zero. All other users are not privileged.

File Privileges

The following configuration files are privileged regardless of their ownership:

- `/etc/ppp/options`
- `/etc/ppp/options.ttyname`
- `/etc/ppp/peers/peer-name`

The file `$HOME/.ppprc` is owned by the user. Options that are read from `$HOME/.ppprc` and from the command line are privileged only if the user who is invoking `pppd` is `root`.

Arguments that follow the `file` option are privileged.

Effects of Option Privileges

Some options require the invoking user or source to be privileged in order to work. Options that are invoked on the command line are assigned the privileges of the user who is running the `pppd` command. These options are not privileged unless the user who is invoking `pppd` is `root`.

Option	Status	Explanation
<code>domain</code>	Privileged	Requires privileges for use.
<code>linkname</code>	Privileged	Requires privileges for use.
<code>noauth</code>	Privileged	Requires privileges for use.
<code>nopam</code>	Privileged	Requires privileges for use.
<code>pam</code>	Privileged	Requires privileges for use.
<code>plugin</code>	Privileged	Requires privileges for use.
<code>privgroup</code>	Privileged	Requires privileges for use.
<code>allow-ip addresses</code>	Privileged	Requires privileges for use.
<code>name hostname</code>	Privileged	Requires privileges for use.
<code>plink</code>	Privileged	Requires privileges for use.
<code>noplink</code>	Privileged	Requires privileges for use.

Composed September 5, 2002

Option	Status	Explanation
plumbed	Privileged	Requires privileges for use.
proxyarp	Becomes privileged if noproxyarp has been specified	Cannot be overridden by an unprivileged user.
defaultroute	Privileged if nodefaultroute is set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
disconnect	Privileged if set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
bsdcomp	Privileged if set in a privileged file or by a privileged user	The non-privileged user cannot specify a code size that is larger than the privileged user has specified.
deflate	Privileged if set in a privileged file or by a privileged user	The non-privileged user cannot specify a code size that is larger than the privileged user has specified.
connect	Privileged if set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
init	Privileged if set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
pty	Privileged if set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
welcome	Privileged if set in a privileged file or by a privileged user	Cannot be overridden by an unprivileged user.
ttyname	Privileged if set in a privileged file	Opened with root permissions regardless of who invokes pppd.
	Not privileged if set in a non-privileged file	Opened with the privileges of the user who invokes pppd.

/etc/ppp/options Configuration File

You use the `/etc/ppp/options` file to define global options for all PPP communications on the local machine. `/etc/ppp/options` is a privileged file. `/etc/ppp/options` should be owned by root, although `pppd` does not enforce this rule. Options that you define in `/etc/ppp/options` have precedence over definitions of the same options in all other files and the command line.

Typical options that you might use in `/etc/ppp/options` include the following:

- **lock**— Enables UUCP-style file locking
- **noauth** — Indicates that the machine does not authenticate callers

Note — The Solaris PPP 4.0 software does not include a default `/etc/ppp/options` file. `pppd` does not require the `/etc/ppp/options` file to work. If a machine does not have an `/etc/ppp/options` file, only root can run `pppd` on that machine.

You must create `/etc/ppp/options` by using a text editor, as shown in “How to Define Communications Over the Serial Line” on page 463. If a machine does not require global options, you can create an empty `/etc/ppp/options` file. Then, both root and regular users can run `pppd` on the local machine.

/etc/ppp/options.tpl Template

The `/etc/ppp/options.tpl` contains helpful comments about the `/etc/ppp/options` file plus three common options for the global `/etc/ppp/options` file.

```
lock
nodefaultroute
noproxyarp
```

Option	Definition
lock	Enables UUCP-style file locking
nodefaultroute	Specifies that no default route is defined
noproxyarp	Disallows proxyarp

To use `/etc/ppp/options.tpl` as the global options file, rename `/etc/ppp/options.tpl` to `/etc/ppp/options`. Then, modify the file contents as needed by your site.

Where to Find Sample `/etc/ppp/options` Files

TABLE 36-2 Examples of the `/etc/ppp/options` File

Example <code>/etc/ppp/options</code>	For Instructions
For a dial-out machine	"How to Define Communications Over the Serial Line" on page 463
For a dial-in server	"How to Define Communications Over the Serial Line (Dial-in Server)" on page 471
For PAP support on a dial-in server	"How to Add PAP Support to the PPP Configuration Files (Dial-in Server)" on page 485
For PAP support on a dial-out machine	"How to Add PAP Support to the PPP Configuration Files (Dial-out Machine)" on page 488
For CHAP support on a dial-in server	"How to Add CHAP Support to the PPP Configuration Files (Dial-in Server)" on page 492

`/etc/ppp/options.ttyname` Configuration File

You can configure the characteristics of communications on the serial line in the `/etc/ppp/options.ttyname` file. `/etc/ppp/options.ttyname` is a privileged file that is read by `pppd` after parsing any existing `/etc/ppp/options` and existing `$HOME/.ppprc` files. Otherwise, `pppd` reads `/etc/ppp/options.ttyname` after parsing `/etc/ppp/options`.

ttyname is used for both dial-up and leased-line links. *ttyname* represents a particular serial port on a machine, such as `cua/a` or `cua/b`, where a modem or ISDN TA might be attached.

When naming the `/etc/ppp/options.ttyname` file, replace the slash (/) in the device name with a dot (.). For example, the options file for device `cua/b` should be named `/etc/ppp/options.cua.b`.

Note – Solaris PPP 4.0 does not require an `/etc/ppp/options.ttyname` file to work correctly. Your server might have only one serial line for PPP. Furthermore, the server requires few options. In this instance, you can specify any required options in another configuration file or on the command line.

Using `/etc/ppp/options.ttyname` on a Dial-in Server

For a dial-up link, you might choose to create individual `/etc/ppp/options.ttyname` files for every serial port on a dial-in server with a modem attached. Typical options include the following:

- IP address required by the dial-in server

Set this option if you require incoming callers on serial port *ttyname* to use a particular IP address. Your address space might have a limited number of IP addresses that are available for PPP in comparison to the number of potential callers. In this situation, consider assigning an IP address to each serial interface that is used for PPP on the dial-in server. This assignment implements dynamic addressing for PPP.

- `asyncmap map_value`

The `asyncmap` option maps control characters that cannot be received over the serial line by the particular modem or ISDN TA. When the `xonxoff` option is used, `pppd` automatically sets an `asyncmap` of `0xa0000`.

map_value states, in hexadecimal format, the control characters that are problematic.

- `init "chat -U -f /etc/ppp/mychat"`

The `init` option tells the modem to initialize communications over the serial line by using the information in the `chat -U` command. The modem uses the chat string in the file `/etc/ppp/mychat`.

- Security parameters that are listed in the `pppd(1m)` man page

Using `/etc/ppp/options.ttyname` on a Dial-out Machine

For a dial-out system, you can create an `/etc/ppp/options.ttyname` file for the serial port with the modem, or choose not to use `/etc/ppp/options.ttyname`.

Note – Solaris PPP 4.0 does not require an `/etc/ppp/options.ttyname` file to work correctly. A dial-out machine might have only one serial line for PPP. Furthermore, the dial-out machine might require few options. You can specify any required options in another configuration file or on the command line.

`options.ttya.tmpl` Template File

The `/etc/ppp/options.ttya.tmpl` file contains helpful comments about the `/etc/ppp/options.tty-name` file. The template contains three common options for the `/etc/ppp/options.tty-name` file.

```
38400
asyncmap 0xa0000
:192.168.1.1
```

Option	Definition
38400	Use this baud rate for port ttya.
asyncmap 0xa0000	Assign the asyncmap value of 0xa0000 so that the local machine can communicate with broken peers.
:192.168.1.1	Assign the IP address 192.168.1.1 to all peers that are calling in over the link.

To use `/etc/ppp/options.ttya.tmpl` at your site, rename `/etc/ppp/options.tmpl` to `/etc/ppp/options.ttya-name`. Replace *ttya-name* with the name of the serial port with the modem. Then modify the file contents as needed by your site.

Where to Find Sample `/etc/ppp/options.ttyname` Files

TABLE 36-3 Examples of the `/etc/ppp/options.ttyname` File

Example <code>/etc/ppp/options.ttyname</code>	For Instructions
For a dial-out machine	"How to Define Communications Over the Serial Line" on page 463
For a dial-in server	"How to Define Communications Over the Serial Line (Dial-in Server)" on page 471

Configuring User-Specific Options

This section contains detailed information on setting up users on the dial-in server.

Configuring `$HOME/.ppprc` on a Dial-in Server

The `$HOME/.ppprc` file is intended for users who are configuring preferred PPP options. As administrator, you can also configure `$HOME/.ppprc` for users.

The options in `$HOME/.ppprc` are privileged only when the user who is invoking the file is privileged.

When a caller uses the `pppd` command to initiate a call, the `.ppprc` file is the second file that is checked by the `pppd` daemon.

See "Setting Up Users of the Dial-in Server" on page 469 for instructions on setting up `$HOME/.ppprc` on the dial-in server.

Configuring `$HOME/.ppprc` on a Dial-out Machine

Note – The `$HOME/.ppprc` is not needed on the dial-out machine for Solaris PPP 4.0 to work correctly.

You do not need to have a `$HOME/.ppprc` on a dial-out machine, except for special circumstances. Create one or more `.ppprc` files if you do the following:

- Allow multiple users with differing communications needs to call remote peers from the same machine. In such an instance, create individual `.ppprc` files in the home directories of each user who must dial out.
- Need to specify options that control problems specific to your link, such as disabling Van Jacobson compression. See James Carlson's *PPP Design, Implementation, and Debugging* and the `pppd(1M)` man page for assistance in troubleshooting link problems.

Because the `.ppprc` file is most often used when configuring a dial-in server, refer to "How to Configure Users of the Dial-in Server" on page 470 for configuration instructions for `.ppprc`.

Specifying Information for Communicating With the Dial-in Server

To communicate with a dial-in server, you need to gather information about the server. Then edit a few files. Most significantly, you must configure the communications requirements of all dial-in servers that the dial-out machine needs to call. You can specify options about a dial-in server, such as an ISP phone number, in the `/etc/ppp/options.ttyname` file. However, the optimum place to configure peer information is in `/etc/ppp/peers/peer-name` files.

`/etc/ppp/peers/peer-name` File

Note – The `/etc/ppp/peers/peer-name` file is not needed on the dial-out machine for Solaris PPP 4.0 to work correctly.

Use the `/etc/ppp/peers/peer-name` file to provide information for communicating with a particular peer. `/etc/ppp/peers/peer-name` allows ordinary users to invoke preselected privileged options that users are not allowed to set.

For example, a non-privileged user cannot override the `noauth` option if `noauth` is specified in the `/etc/ppp/peers/peer-name` file. Suppose the user wants to set up a link to `peerB`, which does not provide authentication credentials. As superuser, you can create a `/etc/ppp/peers/peerB` file that includes the `noauth` option. `noauth` indicates that the local machine does not authenticate calls from `peerB`.

The `pppd` daemon reads `/etc/ppp/peers/peer-name` when `pppd` encounters the following option:

```
call peer-name
```

You can create a `/etc/ppp/peers/peer-name` file for each target peer with which the dial-out machine needs to communicate. This practice is particularly convenient for permitting ordinary users to invoke special dial-out links without needing root privileges.

Typical options that you specify in `/etc/ppp/peers/peer-name` include the following:

- `user user_name`
Supply `user_name` to the dial-in server, as the login name of the dial-out machine, when authenticating with PAP or CHAP.
- `remotename peer-name`
Use `peer-name` as the name of the dial-in machine. `remotename` is used in conjunction with PAP or CHAP authentication when scanning the `/etc/ppp/pap-secrets` or `/etc/ppp/chap-secrets` files.
- `connect "chat chat_script . . ."`
Open communication to the dial-in server, by using the instructions in the chat script.
- `noauth`
Do not authenticate the peer `peer-name` when initiating communications.
- `noipdefault`
Set the initial IP address that is used in negotiating with the peer to 0.0.0.0. Use `noipdefault` when setting up a link to most ISPs to help facilitate IPCP negotiation between the peers.
- `defaultroute`
Install a default IPv4 route when IP is established on the link.

See the `pppd(1M)` man page for more options that might apply to a specific target peer.

/etc/ppp/peers/myisp.tmpl Template File

The `/etc/ppp/peers/myisp.tmpl` file contains helpful comments about the `/etc/ppp/peers/peer-name` file. The template concludes with common options that you might use for an `/etc/ppp/peers/peer-name` file:

```
connect "/usr/bin/chat -f /etc/ppp/myisp-chat"
user myname
remotename myisp
noauth
noipdefault
defaultroute
updetach
noccip
```

Option	Definition
<code>connect "/usr/bin/chat -f /etc/ppp/myisp-chat"</code>	Call the peer by using the chat script <code>/etc/ppp/myisp-chat</code> .
<code>user myname</code>	Use this account name for the local machine. <code>myname</code> is the name for this machine in the peer's <code>/etc/ppp/pap-secrets</code> file.
<code>remotename myisp</code>	Recognize <code>myisp</code> as the name of the peer in the local machine's <code>/etc/ppp/pap-secrets</code> file.
<code>noauth</code>	Do not require calling peers to provide authentication credentials.
<code>noipdefault</code>	Do not use a default IP address for the local machine.
<code>defaultroute</code>	Use the default route that is assigned to the local machine.
<code>updetach</code>	Log errors in the PPP log files, rather than on the standard output.
<code>noccip</code>	Do not use CCP compression.

To use `/etc/ppp/peers/myisp.tmpl` at your site, rename `/etc/ppp/peers/myisp.tmpl` to `/etc/ppp/peers/.peer-name`. Replace `peer-name` with the name of the peer to be called. Then modify the file contents as needed by your site.

Where to Find Sample `/etc/ppp/peers/peer-name` Files

TABLE 36-4 Examples of `/etc/ppp/peers/peer-name` Files

Example <code>/etc/ppp/peers/peer-name</code>	For Instructions
For a dial-out machine	"How to Define the Connection With an Individual Peer" on page 465
For a local machine on a leased line	"How to Configure a Machine on a Leased Line" on page 478
To support PAP authentication on a dial-out machine	"How to Add PAP Support to the PPP Configuration Files (Dial-out Machine)" on page 488
To support CHAP authentication on a dial-out machine	"How to Add CHAP Support to the PPP Configuration Files (Dial-out Machine)" on page 494
To support PPPoE on a client system	"Setting Up the PPPoE Client" on page 498

Configuring Modems for a Dial-up Link

This section contains information about configuring modems.

Configuring the Modem Speed

A major issue in modem configuration is designating the speed at which the modem should operate. The following guidelines apply to modems that are used with Sun Microsystems computers:

- Older SPARC systems – Check the hardware documentation that accompanies the system. Many SPARCstation™ machines require modem speed not to exceed 38400 bps.
- UltraSPARC™ machines – Set the modem speed to 115200 bps, which is useful with modern modems and fast enough for a dial-up link. If you plan to use a dual-channel ISDN TA with compression, you need to increase the modem speed. The limit on an UltraSPARC is 460800 bps for an asynchronous link.

For a *dial-out machine*, set the modem speed in the PPP configuration files, such as `/etc/ppp/peers/peer-name`, or by specifying the speed as an option for `pppd`.

For a *dial-in server*, you need to set the speed by using the `ttymon` facility or `admintool`, as described in “Configuring Devices on the Dial-in Server” on page 468.

Defining the Conversation on the Dial-up Link

The dial-out machine and its remote peer communicate across the PPP link by negotiating and exchanging various instructions. When configuring a dial-out machine, you need to determine what instructions are required by the local and remote modems. Then you create a file that is called a chat script that contains these instructions. This section discusses information about configuring modems and creating chat scripts.

Contents of the Chat Script

Each remote peer that the dial-out machine needs to connect to probably requires its own chat script.

Note – Chat scripts are typically used only on dial-up links. Leased-line links do not use chat scripts unless the link includes an asynchronous interface that requires startup configuration.

The contents of the chat script are determined by the requirements of your modem model or ISDN TA, and the remote peer. These contents appear as a set of expect-send strings. The dial-out machine and its remote peers exchange the strings as part of the communications initiation process.

An *expect* string contains characters that the dial-out host machine expects to receive from the remote peer to initiate conversation. A *send* string contains characters that the dial-out machine sends to the remote peer after receiving the expect string.

Information in the chat script usually includes the following:

- Modem commands, often referred to as *AT commands*, which enable the modem to transmit data over the telephone
- Phone number of the target peer
This phone number might be the number that is required by your ISP, or a dial-in server at a corporate site, or an individual machine.
- Time-out value, if required

- Login sequence that is expected from the remote peer
- Login sequence that is sent by the dial-out machine

Chat Script Examples

This section contains chat scripts that you can use as a reference for creating your own chat scripts. The modem manufacturer's guide and information from your ISP and other target hosts contain chat requirements for the modem and your target peers. In addition, numerous PPP web sites have sample chat scripts.

Basic Modem Chat Script

The following is a basic chat script that you can use as a template for creating your own chat scripts.

```
ABORT    BUSY
ABORT    'NO CARRIER'
REPORT   CONNECT
TIMEOUT  10
" "      AT&F1M0&M5S2=255
SAY      "Calling myserver\n"
TIMEOUT  60
OK       "ATDT1-123-555-1212"
ogin:    pppuser
ssword:  \q\U
%        pppd
```

The next table describes the contents of the chat script.

Script Contents	Explanation
ABORT 'NO CARRIER'	Abort transmission if the modem reports ABORT 'NO CARRIER' when dialing. The cause for this message is usually a dialing or modem negotiation failure.
REPORT CONNECT	Gather the CONNECT string from the modem. Print the string.
TIMEOUT 10	Set initial timeout to 10 seconds. The modem's response should be immediate.
" " AT&F1M0&M5S2=255	M0 – Turn off the speaker during connect. &M5 – Make the modem require error control. S2=255 – Disable the TIES “+++” break sequence.
SAY "Calling myserver\n"	Display the message “Calling myserver” on the local machine.

Script Contents	Explanation
TIMEOUT 60	Reset the timeout to 60 seconds to allow more time for link negotiation.
OK "ATDT1-123-555-1212"	Call the remote peer by using the phone number 123-555-1212.
ogin: pppuser	Log in to the peer by using UNIX-style login. Supply the user name pppuser.
ssword: \q\U	\q – Do not log if debugging with the -v option. \U – Insert in this location the contents of the string that follows -U, which is specified on the command line. Usually, the string contains the password.
% pppd	Wait for the % shell prompt, and run the pppd command.

/etc/ppp/myisp-chat.tpl Chat Script Template

Solaris PPP 4.0 includes the `/etc/ppp/myisp-chat.tpl`, which you can modify for use at your site. `/etc/ppp/myisp-chat.tpl` is similar to the basic modem chat script except that the template does not include a login sequence.

```
ABORT    BUSY
ABORT    'NO CARRIER'
REPORT   CONNECT
TIMEOUT  10
" "      "AT&F1"
OK       "AT&C1&D2"
SAY      "Calling myisp\n"
TIMEOUT  60
OK       "ATDT1-123-555-1212"
CONNECT  \c
```

Script Contents	Explanation
ABORT BUSY	Abort transmission if the modem receives this message from the opposite peer.
ABORT 'NO CARRIER'	Abort transmission if the modem reports ABORT 'NO CARRIER' when dialing. The cause for this message is usually a dialing or modem negotiation failure.
REPORT CONNECT	Gather the CONNECT string from the modem. Print the string.
TIMEOUT 10	Set initial timeout to 10 seconds. The modem's response should be immediate.
" " "AT&F1"	Reset the modem to factory defaults.

Composed September 5, 2002

Script Contents	Explanation
OK "AT&C1&D2"	Reset the modem so that, for &C1, DCD from the modem follows carrier. If the remote side hangs up the phone for some reason, then the DCD drops. For &D2, DTR high-to-low transition causes the modem to go "on-hook" or hang up.
SAY "Calling myisp\n"	Display the message "Calling myisp" on the local machine.
TIMEOUT 60	Reset the timeout to 60 seconds to allow more time for link negotiation.
OK "ATDT1-123-555-1212"	Call the remote peer by using the phone number 123-555-1212.
CONNECT \c	Wait for the CONNECT message from the opposite peer's modem.

Modem Chat Script for Calling an ISP

Use the next chat script as a template for calling an ISP from a dial-out machine with a U.S. Robotics Courier modem.

```
ABORT BUSY
ABORT 'NO CARRIER'
REPORT CONNECT
TIMEOUT 10
"" AT&F1M0&M5S2=255
SAY "Calling myisp\n"
TIMEOUT 60
OK "ATDT1-123-555-1212"
CONNECT \c
\r \d\c
SAY "Connected; running PPP\n"
```

The following table describes the contents of the chat script.

Script Contents	Explanation
ABORT BUSY	Abort transmission if the modem receives this message from the opposite peer.
ABORT 'NO CARRIER'	Abort transmission if the modem receives this message from the opposite peer.
REPORT CONNECT	Gather the CONNECT string from the modem. Print the string.
TIMEOUT 10	Set initial timeout to 10 seconds. The modem's response should be immediate.

Script Contents	Explanation
" " AT&F1M0M0M0M0&M5S2=255	M0 – Turn off the speaker during connect. &M5 – Make the modem require error control. S2=255 – Disable the TIES “+++” break sequence.
SAY "Calling myisp\n"	Display the message “Calling myisp” on the local machine.
TIMEOUT 60	Reset the timeout to 60 seconds to allow more time for link negotiation.
OK "ATDT1-123-555-1212"	Call the remote peer by using the phone number 123-555-1212.
CONNECT \c	Wait for the CONNECT message from the opposite peer’s modem.
\r \d\c	Wait until the end of the CONNECT message.
SAY "Connected; running PPP\n"	Display the informative message "Connected; running PPP" on the local machine.

Basic Chat Script Enhanced for a UNIX-Style Login

The next chat script is a basic script that is enhanced for calling a remote Solaris peer or other UNIX-type peer. This chat script is used in “How to Create the Instructions for Calling a Peer” on page 464.

```

SAY "Calling the peer\n"
TIMEOUT 10
ABORT BUSY
ABORT 'NO CARRIER'
ABORT ERROR
REPORT CONNECT
" " AT&F1&M5S2=255
TIMEOUT 60
OK ATDT1-123-555-1234
CONNECT \c
SAY "Connected; logging in.\n"
TIMEOUT 5
ogin:--ogin: pppuser
TIMEOUT 20
ABORT 'ogin incorrect'
ssword: \qmypassword
"% " \c
SAY "Logged in. Starting PPP on peer system.\n"
ABORT 'not found'
" " "exec pppd"
~ \c

```

The following table explains the parameters of the chat script.

Composed September 5, 2002

Script Contents	Explanation
TIMEOUT 10	Set initial timeout to 10 seconds. The modem's response should be immediate.
ABORT BUSY	Abort transmission if the modem receives this message from the opposite peer.
ABORT 'NO CARRIER'	Abort transmission if the modem receives this message from the opposite peer.
ABORT ERROR	Abort transmission if the modem receives this message from the opposite peer.
REPORT CONNECT	Gather the CONNECT string from the modem. Print the string.
" " AT&F1&M5S2=255	&M5 – Make the modem require error control. S2=255 – Disable the TIES “+++” break sequence.
TIMEOUT 60	Reset the timeout to 60 seconds to allow more time for link negotiation.
OK ATDT1-123-555-1234	Call the remote peer by using the phone number 123-555-1212.
CONNECT \c	Wait for the CONNECT message from the opposite peer's modem.
SAY "Connected; logging in.\n"	Display the informative message “Connected; logging in,” to give the user status.
TIMEOUT 5	Change the timeout to enable quick display of the login prompt.
ogin:--ogin: pppuser	Wait for the login prompt. If the prompt is not received, send a RETURN and wait. Then, send the user name pppuser to the peer. The sequence that follows is referred to by most ISPs as the PAP login. However, the PAP login is not related in any way to PAP authentication.
TIMEOUT 20	Change the timeout to 20 seconds to allow for slow password verification.
ssword: \qmysecrethere	Wait for the password prompt from the peer. When the prompt is received, send the password \qmysecrethere. The \q prevents the password from being written to the system log files.
"% " \c	Wait for a shell prompt from the peer. The chat script uses the C shell. Change this value if the user prefers to log in with a different shell.
SAY "Logged in. Starting PPP on peer system.\n"	Display the informative message “Logged in. Starting PPP on peer system” to give the user status.
ABORT 'not found'	Abort the transmission if the shell encounters errors.
" " "exec pppd"	Start pppd on the peer.
~ \c	Wait for PPP to start on the peer.

Starting PPP right after the `CONNECT \c` is often called a *PAP login* by ISPs, though the PAP login is actually not part of PAP authentication.

The phrase `ogin:--ogin: pppuser` instructs the modem to send the user name `pppuser` in response to the login prompt from the dial-in server. `pppuser` is a special PPP user account name that was created for remote `user1` on the dial-in server. For instructions on creating PPP user accounts on a dial-in server, refer to “How to Configure Users of the Dial-in Server” on page 470.

Chat Script for External ISDN TA

The following chat script is for calling from a dial-out machine with a ZyXEL omni.net. ISDN TA.

```
SAY "Calling the peer\n"
TIMEOUT 10
ABORT BUSY
ABORT 'NO CARRIER'
ABORT ERROR
REPORT CONNECT
"" AT&FB40S83.7=1&K44&J3X7S61.3=1S0=0S2=255
OK ATDI18882638234
CONNECT \c
\r \d\c
SAY "Connected; running PPP\n"
```

The following table explains the parameters of the chat script.

Script Contents	Explanation
SAY "Calling the peer"	Display this message on the screen of the dial-out machine.
TIMEOUT 10	Set the initial timeout to 10 seconds.
ABORT BUSY	Abort transmission if the modem receives this message from the opposite peer.
ABORT 'NO CARRIER'	Abort transmission if the modem receives this message from the opposite peer.
ABORT ERROR	Abort transmission if the modem receives this message from the opposite peer.
REPORT CONNECT	Gather the <code>CONNECT</code> string from the modem. Print the string.

Script Contents	Explanation
""	The letters in this line have the following meaning:
AT&FB40S83.7=	■ &F – Use factory default
1&K44&J3X7S61.3=1	■ B40 – Do asynchronous PPP conversion
S0=0S2=255	■ S83.7=1 – Use data over speech bearer
	■ &K44 – Enable CCP compression
	■ &J3 – Enable MP
	■ X7 – Report DCE side rates
	■ S61.3=1 – Use packet fragmentation
	■ S0=0 – No auto answer
	■ S2=255 – Disable TIES escape
OK ATDI18882638234	Make an ISDN call. For multi-link, the second call is placed to the same telephone number, which is normally what is required by most ISPs. If the remote peer requires a different second phone number, append "+ <i>nnnn</i> ". <i>nnnn</i> represents the second phone number.
CONNECT \c	Wait for the CONNECT message from the opposite peer's modem.
\r \d \c	Wait until the end of the CONNECT message.
SAY "Connected; running PPP\n"	Display this message on the screen of the dial-out machine.

Refer to the `chat(1M)` man page for descriptions of options and other detailed information about the chat script. For an explanation of expect-send strings, refer to "UUCP Chat Script Field" on page 590.

For More Chat Script Examples

A number of web sites offer sample chat scripts and assistance in creating the chat scripts.

The PPP Frequently Asked Questions (FAQ) available from Australian National University posts *URL*.

Invoking the Chat Script

You call chat scripts by using the `connect` option. You can use `connect "chat . . ."` in any PPP configuration file or on the command line.

Chat scripts are not executable, but the program that is invoked by `connect` must be executable. You might use the `chat` utility as the program to be invoked by `connect`. In this instance, if you store the chat script in an external file through the `-f` option, then your chat script file is not executable.

The `chat` program that is described in `chat(1m)` executes the actual chat script. The `pppd` daemon invokes the `chat` program whenever `pppd` encounters the `connect "chat . . . "` option.

Note – You can use any external program, such as Perl or Tcl, to create advanced chat scripts. Solaris PPP 4.0 provides the `chat` utility as a convenience.

▼ How to Invoke a Chat Script (Task)

1. Create the chat script as an ASCII file.
2. Invoke the chat script in any PPP configuration file by using the following syntax:

```
connect 'chat -f /etc/ppp/chatfile'
```

The `-f` flag indicates that a file name is to follow. `/etc/ppp/chatfile` represents the name of the chat file.

3. Give read permission for the external chat file to the user who runs the `pppd` command.



Caution – The chat program always runs with the user's privileges, even if the `connect 'chat . . . '` option is invoked from a privileged source. Thus, a separate chat file that is read with the `-f` option must be readable by the invoking user. This privilege can be a security problem if the chat script contains passwords or other sensitive information.

Chat Script in an External File

If the chat script that is needed for a particular peer is long or complicated, consider creating the script as a separate file. External chat files are easy to maintain and to document. You can add comments to the chat file by preceding the comments with the hash (`#`) sign.

The procedure “How to Create the Instructions for Calling a Peer” on page 464 shows the use of a chat script that is contained in an external file.

Inline Chat Script

You can place the entire chat script conversation on a single line, similar to the following:


```
connect 'chat "" "AT&F1" OK ATDT5551212 CONNECT "\c"'
```

The complete chat script follows the `chat` keyword. The script terminates with `"\c" \`. You use this form in any PPP configuration file or on the command line, as an argument to `pppd`.

Creating a Chat File That Is Executable

You can create a chat file that is an executable script to be run automatically when the dial-up link is initiated. Thus you can run additional commands during link initiation, such as `stty` for parity settings, besides the commands that are contained in a traditional chat script.

This executable chat script logs in to an old-style UNIX system that requires 7 bits with even parity. The system then changes to 8 bits with no parity when running PPP.

```
#!/bin/sh
chat "" "AT&F1" OK "ATDT555-1212" CONNECT "\c"
stty evenp
chat ogin: pppuser ssword: "\q\U" % "exec pppd"
stty -evenp
```

▼ How to Create an Executable Chat Program

1. Use your text editor to create an executable chat program, such as the previous example.
2. Make the chat program executable.

```
# chmod +x /etc/ppp/chatprogram
```

3. Invoke the chat program.

```
connect /etc/ppp/chatprogram
```

Chat programs do not have to be located within the `/etc/ppp` file system. You can store chat programs in any location.

Authenticating Callers on a Link

This section explains how the PPP authentication protocols work and explains the databases that are associated with the authentication protocols.

Password Authentication Protocol (PAP)

PAP authentication is somewhat similar in operation to the UNIX login program, though PAP does not grant shell access to the user. PAP uses the PPP configuration files and PAP database in the form of the `/etc/ppp/pap-secrets` file for setting up authentication. PAP also uses `/etc/ppp/pap-secrets` for defining PAP security credentials. These credentials include a peer name, a “user name” in PAP parlance and a password. PAP credentials also contain related information for each caller who is permitted to link to the local machine. The PAP user names and passwords can be identical to or different from the UNIX user names and passwords in the password database.

`/etc/ppp/pap-secrets` File

The PAP database is implemented in the `/etc/ppp/pap-secrets` file. Machines on both sides of the PPP link must have properly configured PAP credentials in their `/etc/ppp/pap-secrets` files for successful authentication. The caller (authenticatee) supplies credentials in the `user` and `password` columns of the `/etc/ppp/pap-secrets` file or in the obsolete `+ua` file. The server (authenticator) validates these credentials against information in `/etc/ppp/pap-secrets`, through the UNIX `passwd` database, or in the PAM facility.

The `/etc/ppp/pap-secrets` file has the following syntax.

TABLE 36-5 Syntax of `/etc/ppp/pap-secrets`

Caller	Server	Password	IP Addresses
myclient	ISP-server	mypassword	*

The parameters have the following meaning.

myclient	PAP user name of the caller. Often, this name is identical to the caller’s UNIX user name, particularly if the dial-in server uses the <code>login</code> option of PAP.
ISP-server	Name of the remote machine, often a dial-in server.
mypassword	Caller’s PAP password.
IP address	IP address that is associated with the caller. Use an asterisk (*) to indicate any IP address.

Creating PAP Passwords

PAP passwords are sent over the link *in the clear*, that is, in readable ASCII format. For the caller (authenticatee), the PAP password must be stored in the clear in any of the following locations:

Composed September 5, 2002

- In `/etc/ppp/pap-secrets`
- In another external file
- In a named pipe through the `pap-secrets @` feature
- As an option to `pppd`, either on the command line or in a PPP configuration file
- Through the `+ua` file

On the server (authenticator), the PAP password can be hidden by doing one of the following:

- Specifying `papcrypt` and using passwords that are hashed by `crypt(3C)` in the `pap-secrets` file.
- Specifying the `login` option to `pppd` and omitting the password from the `pap-secrets` file by placing double quotes (") in the password column. In this instance, authentication is done through the UNIX `passwd` database or the `pam(3pam)` mechanism.

What Happens During PAP Authentication

PAP authentication occurs in the following sequence.




FIGURE 36-1 PAP Authentication Process

1. The caller (authenticatee) calls the remote peer (authenticator) and provides its PAP user name and password as part of link negotiation.
2. The peer verifies the identity of the caller in its `/etc/ppp/pap-secrets` file. If the peer uses the `login` option of PAP, the peer verifies the caller's user name and password in its password database.
3. If authentication is successful, the peer continues link negotiation with the caller. If authentication fails, the link is dropped.
4. (Optional) If the caller authenticates responses from remote peers, the remote peer must send its own PAP credentials to the caller. Thus, the remote peer becomes the authenticatee and the caller the authenticator.
5. The original caller reads its own `/etc/ppp/pap-secrets` to verify the identity of the remote peer.

Note – If the original caller does require authentication credentials from the remote peer, Step 1 and Step 4 happen in parallel.

If the peer is authenticated, negotiation continues. Otherwise, the link is dropped.

- 6. Negotiation between caller and peer continues until the link is successfully established.

Using the login Option With /etc/ppp/pap-secrets

You can add the login option for authenticating PAP credentials to any PPP configuration file. When login is specified, for example, in /etc/ppp/options, pppd verifies that the caller's PAP credentials exist in the Solaris password database. The following table shows the format of a /etc/ppp/pap-secrets file with the login option.

TABLE 36-6 /etc/ppp/pap-secrets With login Option

Caller	Server	Password	IP Addresses
joe	*	" "	*
sally	*	" "	*
sue	*	" "	*

The parameters have the following meanings.

Caller	Names of all authorized callers.
Server	Asterisk, which indicates that any server name is valid. The name option is not required in the PPP configuration files.
Password	Double quotes, which indicate that any password is valid. If a password is in this column, then the password from the peer must match both the PAP password and the UNIX passwd database.
IP Addresses	Asterisk, which indicates that any IP address is allowed.

Challenge-Handshake Authentication Protocol (CHAP)

CHAP authentication uses the notion of the *challenge* and *response*, which means that the peer (authenticator) challenges the caller (authenticatee) to prove its identity. The challenge includes a random number and a unique ID that is generated by the authenticator. The caller must use the ID, random number, and its CHAP security credentials to generate the proper response (handshake) to send to the peer.

CHAP security credentials include a CHAP user name and a CHAP “secret.” The chat secret is an arbitrary string that is known to both the caller and the peer before they negotiate a PPP link. You configure CHAP security credentials in the CHAP database, `/etc/ppp/chap-secrets`.

`/etc/ppp/chap-secrets` File

The CHAP database is implemented in the `/etc/ppp/chap-secrets` file. Machines on both sides of the PPP link must have each others’ CHAP credentials in their `/etc/ppp/chap-secrets` files for successful authentication.

Note – Unlike PAP, the shared secret must be in the clear on both peers. You cannot use crypt, PAM, or the PPP login option with CHAP.

The `/etc/ppp/chap-secrets` file has the following syntax.

TABLE 36–7 Syntax of `/etc/ppp/chap-secrets`

Caller	Server	CHAP secret	IP Addresses
myclient	myserver	secret5748	*

The parameters have the following meanings:

myclient	CHAP user name of the caller. This name can be the same as or different from the caller’s UNIX user name.
myserver	Name of the remote machine, often a dial-in server.
secret5748	Caller’s CHAP secret. Note – Unlike PAP passwords, CHAP secrets are never sent over the link. Rather, CHAP secrets are used when the local machines compute the response.
IP address	IP address that is associated with the caller. Use an asterisk (*) to indicate any IP address.

What Happens During CHAP Authentication

CHAP authentication occurs in the following sequence.




FIGURE 36-2 CHAP Authentication Sequence

1. Two peers that are about to initiate communications agree on a secret to be used for authentication during negotiation of a PPP link.

2. The administrators of both machines add the secret, CHAP user names, and other CHAP credentials to the `/etc/ppp/chap-secrets` database of their respective machines.
3. The caller (authenticatee) calls the remote peer (authenticator).
4. The authenticator generates a random number and an ID, and sends this data to the authenticatee as a challenge.
5. The authenticatee looks up the peer's name and secret in its `/etc/ppp/chap-secrets` database.
6. The authenticatee calculates a response by applying the MD5 computational algorithm to the secret and the peer's random number challenge. Then the authenticatee sends the results as its response to the authenticator.
7. The authenticator looks up the authenticatee's name and secret in its `/etc/ppp/chap-secrets` database.
8. The authenticator calculates its own figure by applying MD5 to the number that was generated as the challenge and the secret for the authenticatee in `/etc/ppp/chap-secrets`.
9. The authenticator compares its results with the response from the caller. If the two numbers are the same, the peer has successfully authenticated the caller, and link negotiation continues. Otherwise the link is dropped.

Creating an IP Addressing Scheme for Callers

Consider creating one or more IP addresses for all incoming calls instead of assigning a unique IP address to each remote user. Dedicated IP addresses are particularly important if the number of potential callers exceeds the number of serial ports and modems on the dial-in server. You can implement a number of different scenarios, depending on your site's needs. Moreover, the scenarios are not mutually exclusive.

Assigning Dynamic IP Addresses to Callers

Dynamic addressing involves the assignment to each caller of the IP address that is defined in `/etc/ppp/options.ttyname`. Dynamic addressing occurs on a per-serial port basis. When a call arrives over a serial line, the caller receives the IP address in the `/etc/ppp/options.ttyname` file for the call's serial interface.

For example, suppose a dial-in server has four serial interfaces that provide dial-up service to incoming calls:

- For serial port `term/a`, create the file `/etc/ppp/options.term.a` with the following entry:

`:10.1.1.1`

- For serial port `term/b`, create the file `/etc/ppp/options.term.b` with the following entry:

`:10.1.1.2`

- For serial port `term/c`, create the file `/etc/ppp/options.term.c` with the following entry:

`:10.1.1.3`

- For serial port `term/d`, create the file `/etc/ppp/options.term.d` with the following entry:

`:10.1.1.4`

With the previous addressing scheme, an incoming call on serial interface `/dev/term/c` is given the IP address 10.1.1.3 for the duration of the call. After the first caller hangs up, a later call that comes in over serial interface `/dev/term/c` is also given the IP address 10.1.1.3.

The advantages of dynamic addressing include the following:

- You can track PPP network usage down to the serial port.
- You can assign a minimum number of IP addresses for PPP use.
- You can administer IP filtering in a more simplified fashion.

Assigning Static IP Addresses to Callers

If your site implements PPP authentication, you can assign specific, *static* IP addresses to individual callers. In this scenario, every time a dial-out machine calls the dial-in server, the caller receives the same IP address.

You implement static addresses in either the `pap-secrets` or `chap-secrets` database. Here is a sample `/etc/ppp/pap-secrets` file that defines static IP addresses.

Caller	Server	Password	IP Addresses
joe	myserver	joepasswd	10.10.111.240
sally	myserver	sallypasswd	10.10.111.241
sue	myserver	suepasswd	10.10.111.242

Here is a sample `/etc/ppp/chap-secrets` file that defines static IP addresses.

Caller	Server	CHAP secret	IP Addresses
account1	myserver	secret5748	10.10.111.244
account2	myserver	secret91011	10.10.111.245

Assigning IP Addresses by sPPP Unit Number

If you are using either PAP or CHAP authentication, you can assign IP addresses to callers by the sPPP unit number. The next table shows an example of this usage.

Caller	Server	Password	IP Addresses
myclient	ISP-server	mypassword	10.10.111.240/28+

The plus sign (+) indicates that the unit number is added to the IP address. Addresses 10.10.111.240 through 10.10.111.255 are assigned to remote users. sPPP0 gets IP address 10.10.111.240. sPPP1 gets IP address 10.10.111.241, and so on.

Creating PPPoE Tunnels for DSL Support

By using PPPoE, you can provide PPP over high-speed digital services to multiple clients that are using one or more DSL modems. PPPoE implements these services by creating an Ethernet tunnel through three participants, the enterprise, the telephone company, and the service provider.

- For an overview and description of how PPPoE works, see “PPPoE Overview” on page 437.
- For tasks for setting up PPPoE tunnels, see Chapter 34.

This section contains detailed information about PPPoE commands and files, which is summarized in the next table.

TABLE 36-8 PPPoE Commands and Configuration Files

File or Command	Description	For Instructions
<code>/etc/ppp/pppoe</code>	A file that contains characteristics that are applied by default to all tunnels that were set up by PPPoE on the system	" <code>/etc/ppp/pppoe</code> File" on page 558
<code>/etc/ppp/pppoe.device</code>	A file that contains characteristics of a particular interface that is used by PPPoE for a tunnel	" <code>/etc/ppp/pppoe.device</code> File" on page 559
<code>/etc/ppp/pppoe.if</code>	File that lists the Ethernet interface over which the tunnel that is set up by PPPoE runs	" <code>/etc/ppp/pppoe.if</code> File" on page 555
<code>/usr/sbin/sppptun</code>	Command for configuring the Ethernet interfaces that are involved in a PPPoE tunnel	" <code>/usr/sbin/sppptun</code> Command" on page 556
<code>/usr/lib/inet/pppoed</code>	Command and options for using PPPoE to set up a tunnel	" <code>/usr/lib/inet/pppoed</code> Daemon" on page 557

Files for Configuring Interfaces for PPPoE

The interfaces that are used at either end of the PPPoE tunnel must be configured before the tunnel can support PPP communications. Use `/usr/sbin/sppptun` and `/etc/ppp/pppoe.if` files for this purpose. You must use these tools to configure Ethernet interfaces on all Solaris PPPoE clients and PPPoE access servers.

`/etc/ppp/pppoe.if` File

The `/etc/ppp/pppoe.if` file lists the names of all Ethernet interfaces on a host to be used for the PPPoE tunnels. This file is processed during system boot when the interfaces that are listed are plumbed for use in PPPoE tunnels.

You need to explicitly create `/etc/ppp/pppoe.if`. Type the name of one interface to be configured for PPPoE on each line.

Sample `/etc/ppp/pppoe.if` File

The following example shows an `/etc/ppp/pppoe.if` file for a server that offers three interfaces for PPPoE tunnels.

```
# cat /etc/ppp/pppoe.if
hme1
hme2
hme3
```

PPPoE clients usually have only one interface that is listed in `/etc/ppp/pppoe.if`.

`/usr/sbin/sppptun` Command

You can use the `/usr/sbin/sppptun` command to manually plumb and unplumb the Ethernet interfaces to be used for PPPoE tunnels. By contrast, `/etc/ppp/pppoe.if` is only read when the system boots. These interfaces should correspond to the interfaces that are listed in `/etc/ppp/pppoe.if`.

`sppptun` plumbs the Ethernet interfaces that are used in PPPoE tunnels in a manner that is similar to the `ifconfig` command. Unlike `ifconfig`, you must plumb interfaces twice to support PPPoE because two Ethernet protocol numbers are involved.

The basic syntax for `sppptun` is as follows:

```
# /usr/sbin/sppptun plumb pppoed device-name
device-name:pppoed
# /usr/sbin/sppptun plumb pppoe device-name
device-name:pppoe
```

In this syntax, *device-name* is the name of the device to be plumbed for PPPoE.

The first time that you issue the `sppptun` command, the discovery protocol `pppoed` is plumbed on the interface. The second time that you run `sppptun`, the session protocol `pppoe` is plumbed. `sppptun` prints the name of the interface that was just plumbed. You use this name to unplumb the interface, when necessary.

For more information, refer to the `sppptun(1M)` man page.

Sample sppptun Commands for Administering Interfaces

- The following sample shows how to manually plumb an interface for PPPoE by using `/usr/sbin/sppptun`.

EXAMPLE 36-1 To plumb an Interface to Support PPPoE

```
# /usr/sbin/sppptun plumb pppoed hme0
hme0:pppoed
# /dev/sppptun plumb pppoe hme0
hme0:pppoe
```

- This sample shows how to list the interfaces on an access server that was plumbed for PPPoE.

EXAMPLE 36-2 To List All Interfaces on a PPPoE Access Server

```
/usr/sbin/sppptun query
hme0:pppoe
hme0:pppoed
```

EXAMPLE 36-2 To List All Interfaces on a PPPoE Access Server (Continued)

```
hme1:pppoe
hme1:pppoed
hme2:pppoe
hme2:pppoed
```

- This sample shows how to unplumb an interface.

EXAMPLE 36-3 To Unplumb an Interface With a PPPoE Tunnel

```
# sppptun unplumb hme0:pppoed
# sppptun unplumb hme0:pppoe
```

PPPoE Access Server Commands and Files

A service provider that offers DSL services or support to customers can use an access server that is running Solaris PPPoE. The PPPoE access server and client do function in the traditional client-server relationship. This relationship is similar to the relationship of the dial-out machine and dial-in server on a dial-up link. One PPPoE system initiates communications and one PPPoE system answers. By contrast, the PPP protocol has no notion of the client-server relationship. PPP considers both systems equal peers.

The commands and files that set up a PPPoE access server include the following:

- “/usr/sbin/sppptun Command” on page 556
- “/usr/lib/inet/pppoed Daemon” on page 557
- “/etc/ppp/pppoe File” on page 558
- “/etc/ppp/pppoe.device File” on page 559
- “pppoe.so Plugin” on page 562

/usr/lib/inet/pppoed Daemon

The pppoed daemon accepts broadcasts for services from prospective PPPoE clients. Additionally, pppoed negotiates the server side of the PPPoE tunnel and runs pppd, the PPP daemon, over that tunnel.

You configure pppoed services in the /etc/ppp/pppoe and /etc/ppp/pppoe.device files. If /etc/ppp/pppoe exists when the system boots, pppoed runs automatically. You can also explicitly run the pppoed daemon on the command line by typing /usr/lib/inet/pppoed.

/etc/ppp/pppoe File

The `/etc/ppp/pppoe` file describes the services that are offered by an access server plus options that define how PPP runs over the PPPoE tunnel. You can define services for individual interfaces, or globally, that is, for all interfaces on the access server. The access server sends the information in the `/etc/ppp/pppoe` file in response to a broadcast from a potential PPPoE client.

The following is the basic syntax of `/etc/ppp/pppoe`:

```
global-options
service service-name
    service-specific-options
    device interface-name
```

The parameters have the following meanings.

<i>global-options</i>	Sets the default options for the <code>/etc/ppp/pppoe</code> file. These options can be any options that are available through <code>pppoed</code> or <code>pppd</code>. For complete lists of options, see the man pages <code>pppoed(1M)</code> and <code>pppd(1M)</code>. For example, you must list the Ethernet interfaces that are available for the PPPoE tunnel as part of <i>global options</i>. If you do not define devices in <code>/etc/ppp/pppoe</code>, the services are not offered on any interface. To define devices as a global option, use the following form: <pre>device interface <,interface></pre> <i>interface</i> specifies the interface where the service listens for potential PPPoE clients. If more than one interface is associated with the service, separate each name with a comma.
<i>service service-name</i>	Starts the definition of the service <i>service-name</i>. <i>service-name</i> is a string that can be any phrase that is appropriate to the services that are provided.
<i>service-specific-options</i>	Lists the PPPoE and PPP options specific to this service.
<i>device interface-name</i>	Specifies the interface where the previously listed service is available.

For additional options to `/etc/ppp/pppoe`, refer to the `pppoed(1M)` and `pppd(1M)` man pages.

A typical `/etc/ppp/pppoe` file might resemble the following.

EXAMPLE 36-4 Basic `/etc/ppp/pppoe` File

```
device hme1,hme2,hme3
service internet
    pppd "name internet-server"
service intranet
    pppd "192.168.1.1:"
```

EXAMPLE 36-4 Basic `/etc/ppp/pppoe` File (Continued)

```

service debug
    device hme1
    pppd "debug name internet-server"

```

In this file, the following values apply.

<code>hme1, hme2, hme3</code>	Three interfaces on the access server to be used for PPPoE tunnels.
<code>service internet</code>	Advertises a service that is called <code>internet</code> to prospective clients. The provider that offers the service also determines how <code>internet</code> is defined. For example, a provider might interpret <code>internet</code> to mean various IP services, as well as access to the Internet.
<code>pppd</code>	Sets the command-line options that are used when the caller invokes <code>pppd</code> . The option <code>"name internet-server"</code> gives the name of the local machine, the access server, as <code>internet-server</code> .
<code>service intranet</code>	Advertises another service that is called <code>intranet</code> to prospective clients.
<code>pppd "192.168.1.1:"</code>	Sets the command-line options that are used when the caller invokes <code>pppd</code> . When the caller invokes <code>pppd</code> , <code>192.168.1.1</code> is set as the IP address for the local machine, the access server.
<code>service debug</code>	Advertises a third service, debugging, on the interfaces that are defined for PPPoE.
<code>device hme1</code>	Restricts debugging to PPPoE tunnels to <code>hme1</code> .
<code>pppd "debug name internet-server"</code>	Sets the command-line options that are used when the caller invokes <code>pppd</code> , in this instance, PPP debugging on <code>internet-server</code> , the local machine.

`/etc/ppp/pppoe.device` File

The `/etc/ppp/pppoe.device` file describes the services that are offered on one interface of a PPPoE access server. `/etc/ppp/pppoe.device` also includes options that define how PPP runs over the PPPoE tunnel. `/etc/ppp/pppoe.device` is an optional file, which operates exactly like the global `/etc/ppp/pppoe`. However, if `/etc/ppp/pppoe.device` is defined for an interface, its parameters have precedence for that interface over the global parameters that are defined in `/etc/ppp/pppoe`.

The basic syntax of `/etc/ppp/pppoe.device` is as follows:

```

service service-name
    service-specific-options
service another-service-name
    service-specific-options

```

The only difference between this syntax and the syntax of `/etc/ppp/pppoe` is that you cannot use the device option that is shown in “`/etc/ppp/pppoe` File” on page 558.

pppoe . so Plugin

`pppoe . so` is the PPPoE shared object file that must be invoked by PPPoE access servers and clients. This file limits MTU and MRU to 1492, filters packets from the driver, and negotiates the PPPoE tunnel, along with `pppoed`. On the access server side, `pppoe . so` is automatically invoked by the `pppd` daemon.

Using PPPoE and PPP Files to Configure an Access Server

This section contains samples of all files that are used to configure an access server. The access server is multihomed. The server is attached to three subnets: green, orange, and purple. `pppoed` runs as root on the server, which is the default.

PPPoE clients can access the orange and purple networks through interfaces `hme0` and `hme1`. Clients log in to the server by using the standard UNIX login. The server authenticates the clients by using PAP.

The green network is not advertised to clients. The only way clients can access green is by directly specifying “green-net” and supplying CHAP authentication credentials. Moreover, only clients `joe` and `mary` are allowed to access the green network by using static IP addresses.

EXAMPLE 36-5 `/etc/ppp/pppoe` File for an Access Server

```
service orange-net
    device hme0,hme1
    pppd "require-pap login name orange-server orange-server:"
service purple-net
    device hme0,hme1
    pppd "require-pap login name purple-server purple-server:"
service green-net
    device hme1
    pppd "require-chap name green-server green-server:"
nowildcard
```

This sample describes the services that are available from the access server. The first service section describes the services of the orange network.

```
service orange-net
    device hme0,hme1
    pppd "require-pap login name orange-server orange-server:"
```

Clients access the orange network over interfaces `hme0` and `hme1`. The options that are given to the `pppd` command force the server to require PAP credentials from potential clients. The `pppd` options also set the server’s name to `orange-server`, as used in the `pap-secrets` file.

The service section for the purple network is identical to the service section of the orange network except for the network and server names.

The next section describes the services of the green network:

```
service green-net
    device hme1
    pppd "require-chap name green-server green-server:"
    nowildcard
```

This section restricts client access to interface hme1. Options that are given to the pppd command force the server to require CHAP credentials from prospective clients. The pppd options also set the server name to green-server, to be used in the chap-secrets file. The nowildcard option specifies that the existence of the green network is not advertised to clients.

For this access server scenario just discussed, you might set up the following /etc/ppp/options file.

EXAMPLE 36-6 /etc/ppp/options File for an Access Server

```
auth
proxyarp
nodefaultroute
name no-service      # don't authenticate otherwise
```

The option name no-service overrides the server name that is normally searched for during PAP or CHAP authentication. The server's default name is found in the /usr/bin/hostname file. The name option in the previous example changes the server's name to no-service. The name no-service is not likely to be found in a pap or chap-secrets file. This action prevents a random user from running pppd and overriding the auth and name options that are set in /etc/ppp/options. pppd then fails because no secrets can be found for the client with a server name of no-service.

The access server scenario uses the following /etc/hosts file.

EXAMPLE 36-7 /etc/hosts File for an Access Server

```
172.16.0.1    orange-server
172.17.0.1    purple-server
172.18.0.1    green-server
172.18.0.2    joes-pc
172.18.0.3    marys-pc
```

Here is the /etc/ppp/pap-secrets file that is used for PAP authentication for clients that attempt to access the orange and purple networks.

EXAMPLE 36-8 /etc/ppp/pap-secrets File for an Access Server

```
* orange-server "" 172.16.0.2/16+
* purple-server "" 172.17.0.2/16+
```

Here is the `/etc/ppp/chap-secrets` file that is used for CHAP authentication. Note that only clients `joe` and `mary` are listed in the file.

EXAMPLE 36-9 `/etc/ppp/chap-secrets` File for an Access Server

```
joe green-server "joe's secret" joes-pc
mary green-server "mary's secret" marys-pc
```

PPPoE Client Commands and Files

To run PPP over a DSL modem, a machine must become a PPPoE client. You have to plumb an interface to run PPPoE, and then use the `pppoec` utility to “discover” the existence of an access server. Thereafter, the client can create the PPPoE tunnel over the DSL modem and run PPP.

The PPPoE client relates to the access server in the traditional client-server model. The PPPoE tunnel is not a dial-up link, but the tunnel is configured and operated in much the same manner.

The commands and files that set up a PPPoE client include the following:

- “`/usr/sbin/spptun` Command” on page 556
- “`/usr/lib/inet/pppoec` Utility” on page 562
- “`pppoe.so` Plugin” on page 562
- “`/etc/ppp/peers/peer-name` File” on page 532
- “`/etc/ppp/options` Configuration File” on page 528

`/usr/lib/inet/pppoec` Utility

The `/usr/lib/inet/pppoec` utility is responsible for negotiating the client side of a PPPoE tunnel. `pppoec` is similar to the Solaris PPP 4.0 `chat` utility. You do not invoke `pppoec` directly. Rather, you start `/usr/lib/inet/pppoec` as an argument to the `connect` option of `pppd`.

`pppoe.so` Plugin

`pppoe.so` is the PPPoE shared object that must be loaded by PPPoE to provide PPPoE capability to access servers and clients. The `pppoe.so` shared object limits MTU and MRU to 1492, filters packets from the driver, and handles runtime PPPoE messages.

On the client side, `pppd` loads `pppoe.so` when the user specifies the `plugin pppoe.so` option.

`/etc/ppp/peers/peer-name` File for Defining an Access Server Peer

When you define an access server to be discovered by `pppoe`, you use options that apply to both `pppoe` and the `pppd` daemon. A `/etc/ppp/peers/peer-name` file for an access server requires the following parameters:

- `sppptun` – Name for the serial device that is used by the PPPoE tunnel
- `plugin pppoe.so` – Instructs `pppd` to load the `pppoe.so` shared object
- `connect "/usr/lib/inet/pppoe device"` – Starts a connection. `connect` then invokes the `pppoe` utility over *device*, the interface that is plumbed for PPPoE

The remaining parameters in the `/etc/ppp/peers/peer-name` file should apply to the PPP link on the server. Use the same options that you would for `/etc/ppp/peers/peer-name` on a dial-out machine. Try to limit the number of options to the minimum you need for the PPP link.

The following example is introduced in “How to Define a PPPoE Access Server Peer” on page 499.

EXAMPLE 36–10 `/etc/ppp/peers/peer-name` to Define a Remote Access Server

```
# vi /etc/ppp/peers/dslserve
sppptun
plugin pppoe.so
connect "/usr/lib/inet/pppoe hme0"
noccp
noauth
user Red
password redsecret
noipdefault
defaultroute
```

This file defines parameters to be used when setting up a PPPoE tunnel and PPP link to access server `dslserve`. The options that are included are as follows.

Option	Description
<code>sppptun</code>	Defines <code>sppptun</code> as the name of the serial device.
<code>plugin pppoe.so</code>	Instructs <code>pppd</code> to load the <code>pppoe.so</code> shared object.
<code>connect "/usr/lib/inet/pppoe hme0"</code>	Runs <code>pppoe</code> and designates <code>hme0</code> as the interface for the PPPoE tunnel and PPP link.

Composed September 5, 2002

Option	Description
noccp	Turns off CCP compression on the link. Note – Many ISPs use only proprietary compression algorithms. Turning off the publicly available CCP algorithm saves negotiation time and avoids very occasional interoperability problems.
noauth	Stops pppd from demanding authentication credentials from the access server. Most ISPs do not provide authentication credentials to customers.
user Red	Sets the name Red as the user name for the client, which is required for PAP authentication by the access server.
password redsecret	Defines redsecret as the password to be provided to the access server for PAP authentication.
noipdefault	Assigns 0.0.0.0 as the initial IP address.
defaultroute	Tells pppd to install a default IPv4 route after IPCP negotiation. You should include defaultroute in <i>/etc/ppp/peers/peer-name</i> when the link is the system's link to the Internet, which is true for a PPPoE client.

Migrating From Asynchronous Solaris PPP to Solaris PPP 4.0 (Tasks)

Earlier versions of the Solaris operating system included a different PPP implementation, Asynchronous Solaris PPP (asppp). If you want to convert peers that run asppp to the newer PPP 4.0, you need to run a conversion script. This chapter covers the following topics in PPP conversion:

- “Before Converting asppp Files ” on page 565
- “Running the asppp2pppd Conversion Script (Tasks)” on page 568

The chapter uses a sample asppp configuration to explain how to accomplish PPP conversion. For a description of the differences between Solaris PPP 4.0 and asppp, go to “Which Version of Solaris PPP to Use” on page 426.

Before Converting asppp Files

You can use the conversion script `/usr/sbin/asppp2pppd` to convert the files that compose a standard asppp configuration:

- `/etc/asppp.cf` – Asynchronous PPP configuration file
- `/etc/uucp/Systems` – UUCP file that describes the characteristics of the remote peer
- `/etc/uucp/Devices` – UUCP file that describes the modem on the local machine
- `/etc/uucp/Dialers` – UUCP file that contains the login sequence to be used by the modem that is described in the `/etc/uucp/Devices` file

For more information about asppp, see the *Solaris 8 System Administration Collection, Volume 3*, available from <http://www.docs.sun.com>.

Example—/etc/asppp.cf Configuration File

The procedure that is shown in “How to Convert From asppp to Solaris PPP 4.0” on page 568 uses the following /etc/asppp.cf file.

```
#
ifconfig ipdptp0 plumb mojave gobi up

path
  inactivity_timeout 120      # Approx. 2 minutes
  interface ipdptp0
  peer_system_name Pgobi      # The name we log in with (also in
                              # /etc/uucp/Systems
```

The file contains the following parameters.

ifconfig ipdptp0 plumb mojave gobi up	Runs the ifconfig command to configure a link from PPP interface ipdptp0 on the local machine mojave to the remote peer gobi
inactivity_timeout 120	Terminates the line after two minutes of inactivity
interface ipdptp0	Configures the interface ipdptp0 on the dial-out machine for asynchronous PPP
peer_system_name Pgobi	Gives the name of the remote peer, Pgobi

Example—/etc/uucp/Systems File

The procedure that is shown in “How to Convert From asppp to Solaris PPP 4.0” on page 568 uses the following /etc/uucp/Systems file.

```
#ident  "@(#)Systems    1.5  92/07/14  SMI"  /* from SVR4 bnu:Systems 2.4 */
#

# .
# .
Pgobi Any ACU 38400 15551212 in:--in: mojave word: sand
```

The file contains the following parameters:

Pgobi	Uses Pgobi as the host name of the remote peer.
Any ACU	Tells the modem on the dial-out machine mojave to establish a link with a modem on Pgobi at any time of the day. Any ACU means “look for ACU in the /etc/uucp/Devices file.”

38400	Sets 38400 as the maximum speed of the link.
15551212	Gives the telephone number of Pgobi.
in:-in: mojave word: sand	Defines the login script that is required by Pgobi to authenticate dial-out machine mojave.

Example—/etc/uucp/Devices File

The procedure that is shown in “How to Convert From asppp to Solaris PPP 4.0” on page 568 uses the following /etc/uucp/Devices file.

```
#ident  "@(#)Devices  1.6  92/07/14  SMI"  /* from SVR4 bnu:Devices 2.7 */

.
.
#

TCP,et - - Any TCP -
.
.
#
ACU cua/b - Any hayes
# 0-7 are on a Magma 8 port card
Direct cua/0 - Any direct
Direct cua/1 - Any direct
Direct cua/2 - Any direct
Direct cua/3 - Any direct
Direct cua/4 - Any direct
Direct cua/5 - Any direct
Direct cua/6 - Any direct
Direct cua/7 - Any direct
# a is the console port (aka "tip" line)
Direct cua/a - Any direct
# b is the aux port on the motherboard
Direct cua/b - Any direct
# c and d are high speed sync/async ports
Direct cua/c - Any direct
Direct cua/d - Any direct
```

The file supports any Hayes modem that is connected to serial port cua/b.

Example—/etc/uucp/Dialers File

The procedure that is shown in “How to Convert From asppp to Solaris PPP 4.0” on page 568 uses the following /etc/uucp/Dialers file.

```
#
#<Much information about modems supported by Solaris UUCP>
```

```
penril      =W-P      "" \d > Q\c : \d- > s\p9\c )-W\p\r\ds\p9\c-) y\c : \E\TP > 9\c OK
ventel      =&-%      "" \r\p\r\c $ k\c ONLINE!
vadic       =K-K      "" \005\p *- \005\p- * \005\p- * D\p BER? \E\T\e \r\c LINE
develcon    ""        "" \pr\ps\c est:\007 \E\D\e \n\007
micom       ""        "" \s\c NAME? \D\r\c GO
direct
#
#
#
# Hayes Smartmodem -- modem should be set with the configuration
# switches as follows:
#
#       S1 - UP          S2 - UP          S3 - DOWN      S4 - UP
#       S5 - UP          S6 - DOWN        S7 - ?          S8 - DOWN
#
hayes       =, -,      "" \dA\pTElVlXlQ0S2=255S12=255\r\c OK\r \EATDT\T\r\c CONNECT
```

<much more information about modems supported by Solaris UUCP>

This file contains the chat scripts for all types of modems, including the Hayes modems that are supported in the /etc/UUCP/Dialers file.

Running the asppp2pppd Conversion Script (Tasks)

The /usr/sbin/asppp2pppd script copies the PPP information in /etc/asppp.cf and PPP-related UUCP files to appropriate locations in the Solaris PPP 4.0 files.

Prerequisites

Before doing the next task, you should have done the following:

- Installed the Solaris 9 operating environment on the machine that also has the asppp and UUCP configuration files
- Become superuser on the machine with the PPP files, for example, the machine `mojave`

▼ How to Convert From asppp to Solaris PPP 4.0

1. Start the conversion script.

```
# /usr/sbin/asppp2pppd
```

The conversion process starts and gives you the following screen output.

This script provides only a suggested translation for your existing aspppd configuration. You will need to evaluate for yourself whether the translation is appropriate for your operating environment.

Continue [Yn]?

2. Type Y to continue. You receive the following output.

```
Chat cannot do echo checking; requests for this removed.
Adding 'noauth' to /etc/ppp/options
```

```
Preparing to write out translated configuration:
```

```
1 chat file:
  1. /etc/ppp/chat.Pgobi.hayes
2 option files:
  2. /etc/ppp/peers/Pgobi
  3. /etc/ppp/options
1 script file:
  4. /etc/ppp/demand
```

The new Solaris PPP 4.0 files have been generated.

▼ How to View the Results of the Conversion

You can view the Solaris PPP 4.0 files that were created by the /usr/sbin/asppp2pppd conversion script at the end of the conversion process. The script displays the following list of options.

```
Enter option number:
```

```
1 - view contents of file on standard output
2 - view contents of file using /usr/bin/less
3 - edit contents of file using /usr/bin/vi
4 - delete/undelete file from list
5 - rename file in list
6 - show file list again
7 - escape to shell (or "!")
8 - abort without saving anything
9 - save all files and exit (default)
```

```
Option:
```

1. Type 1 to view the contents of the files on the screen.

The script requests the number of the file you want to view.

```
File number (1 .. 4):
```

The numbers refer to the translated files that are listed during the conversion process, as shown in the previous Step 2.

2. Type 1 to view the chat file /etc/ppp/chat.Pgobi.hayes.

```
File number (1 .. 4): 1
"" \d\dA\p\pTE1V1X1Q0S2=255S12=255\r\c
OK\r ATDT\T\r\c
CONNECT \c
```

```
in:--in: mojave
word: sand
```

The chat script contains the modem “chat” information that appears on the hayes line in the sample `/etc/UUCP/Dialers` file. `/etc/ppp/chat.Pgobi.hayes` also contains the login sequence for Pgobi that appears in the sample `/etc/UUCP/Systems` file. The chat script is now in the `/etc/ppp/chat.Pgobi.hayes` file.

3. Type 2 to view the peers file, `/etc/ppp/peers/Pgobi`.

```
File number (1 .. 4): 2
/dev/cua/b
38400
demand
idle 120
connect "/usr/bin/chat -f /etc/ppp/chat.Pgobi.hayes -T '15551212'"
user NeverAuthenticate
mojave:gobi
```

The serial port information (`/dev/cua/b`) is from the `/etc/UUCP/Devices` file. The link speed, idle time, authentication information, and peer names are from the `/etc/asppp.cf` file. “demand” refers to the “demand” script, to be called when the dial-out machine tries to connect to peer Pgobi.

4. Type 3 to view the `/etc/ppp/options` file that are created for dial-out machine mohave.

```
File number (1 .. 4): 3
#lock
noauth
```

The information in `/etc/ppp/options` is from the `/etc/asppp.cf` file.

5. Type 4 to view the contents of the demand script.

```
File number (1 .. 4): 4
/usr/bin/pppd file /etc/ppp/peers/Pgobi
```

This script, when invoked, runs the `pppd` command, which then reads the `/etc/ppp/peers/Pgobi` to initiate the link between mohave and Pgobi.

6. Type 9 to save the created files. Then exit the conversion script.

Overview of UUCP

This chapter introduces the UNIX-to-UNIX Copy Program (UUCP) and daemons. The following topics are covered:

- “UUCP Hardware Configurations” on page 571
- “UUCP Software” on page 572
- “UUCP Database Files” on page 574

UUCP enables computers to transfer files and exchange mail with each other. The program also enables computers to participate in large networks such as Usenet.

The Solaris environment provides the Basic Network Utilities (BNU) version of UUCP, also known as HoneyDanBer UUCP. The term *UUCP* denotes the complete range of files and utilities that compose the system, of which the program `uucp` is only a part. The UUCP utilities range from those utilities that are used to copy files between computers (`uucp` and `uuto`) to those utilities that are used for remote login and command execution (`cu` and `uux`).

UUCP Hardware Configurations

UUCP supports the following hardware configurations:

Direct links	You can create a direct link to another computer by running RS-232 cables between serial ports on the two machines. Direct links are useful when two computers communicate regularly and are physically close—within 50 feet of each other. You can use a limited-distance modem to increase this distance somewhat.

Telephone lines	Using an automatic call unit (ACU), such as a high-speed modem, your machine can communicate with other computers over standard phone lines. The modem dials the telephone number that is requested by UUCP. The recipient machine must have a modem capable of answering incoming calls.
Network	UUCP can also communicate over a network that runs TCP/IP or an other protocol family. After your computer has been established as a host on a network, it can contact any other host that is connected to the network.

This chapter assumes that your UUCP hardware has already been assembled and configured. If you need to set up a modem, refer to *System Administration Guide: Basic Administration* and the manuals that accompanied the modem for assistance.

UUCP Software

The UUCP software is automatically included when you run the Solaris installation program and select the entire distribution. Alternatively, you can add it by using `pkgadd`. The UUCP programs can be divided into three categories: daemons, administrative programs, and user programs.

UUCP Daemons

The UUCP system has four daemons: `uucico`, `uuxqt`, `uusched`, and `in.uucpd`. These daemons handle UUCP file transfers and command executions. You can also run them manually from the shell, if necessary.

<code>uucico</code>	Selects the device that is used for the link, establishes the link to the remote computer, and performs the required login sequence and permission checks. Also, <code>uucico</code> transfers data and execute files, logs results, and notifies the user by mail of transfer completions. <code>uucico</code> acts as the “login shell” for UUCP login accounts. When the local <code>uucico</code> daemon calls a remote machine, it communicates directly with the remote <code>uucico</code> daemon during the session. After all the required files have been created, <code>uucp</code>, <code>uuto</code>, and <code>uux</code> programs execute the <code>uucico</code> daemon to contact the remote computer. <code>uusched</code> and <code>Uutry</code> all execute <code>uucico</code>. See the <code>uucico(1M)</code> man page for details.

uuxqt	Executes remote execution requests. This daemon searches the spool directory for execute files (always named <i>X.file</i>) that have been sent from a remote computer. When an <i>X.file</i> file is found, uuxqt opens it to get the list of data files that are required for the execution. uuxqt then checks to see if the required data files are available and accessible. If the files are available, uuxqt checks the <code>Permissions</code> file to verify that it has permission to execute the requested command. The uuxqt daemon is executed by the <code>uudemon.hour</code> shell script, which is started by <code>cron</code> . See the <code>uuxqt(1M)</code> man page for details.
uusched	Schedules the queued work in the spool directory. uusched is initially run at boot time by the <code>uudemon.hour</code> shell script, which is started by <code>cron</code> . See the <code>uusched(1M)</code> man page for details. Before starting the <code>uucico</code> daemon, uusched randomizes the order in which remote computers are called.
in.uucpd	Supports UUCP connections over networks. The <code>inetd</code> on the remote host invokes <code>in.uucpd</code> whenever a UUCP connection is established. uucpd then prompts for a login name. <code>uucico</code> on the calling host must respond with a login name. <code>in.uucpd</code> then prompts for a password, unless one is not required. See the <code>in.uucpd(1M)</code> man page for details.

UUCP Administrative Programs

Most UUCP administrative programs are in `/usr/lib/uucp`. Most basic database files are in `/etc/uucp`. The only exception is `uulog`, which is in `/usr/bin`. The home directory of the `uucp` login ID is `/usr/lib/uucp`. When running the administrative programs through `su` or `login`, use the `uucp` user ID. The user ID owns the programs and spooled data files.

uulog	Displays the contents of a specified computer's log files. Log files are created for each remote computer with which your machine communicates. The log files record each use of <code>uucp</code> , <code>uuto</code> , and <code>uux</code> . See the <code>uucp(1C)</code> man page for details.
uucleanup	Cleans up the spool directory. uucleanup is normally executed from the <code>uudemon.cleanup</code> shell script, which is started by <code>cron</code> . See the <code>uucleanup(1M)</code> man page for details.
Uutry	Tests call-processing capabilities and does moderate debugging. Uutry invokes the <code>uucico</code> daemon to establish a communication link between your machine and the remote computer you specify. See the <code>Uutry(1M)</code> man page for details.
uucheck	Checks for the presence of UUCP directories, programs, and support files. uucheck can also check certain parts of the <code>/etc/uucp/Permissions</code> file for obvious syntactic errors. See the <code>uucheck(1M)</code> man page for details.

UUCP User Programs

The UUCP user programs are in `/usr/bin`. You do not need special permission to use these programs.

<code>cu</code>	Connects your machine to a remote computer so that you can log in to both machines at the same time. <code>cu</code> enables you to transfer files or execute commands on either machine without dropping the initial link. See the <code>cu(1C)</code> man page for details.
<code>uucp</code>	Lets you copy a file from one machine to another. <code>uucp</code> creates work files and data files, queues the job for transfer, and calls the <code>uucico</code> daemon, which in turn attempts to contact the remote computer. See the <code>uucp(1C)</code> man page for details.
<code>uuto</code>	Copies files from the local machine to the public spool directory <code>/var/spool/uucppublic/receive</code> on the remote machine. Unlike <code>uucp</code> , which lets you copy a file to any accessible directory on the remote machine, <code>uuto</code> places the file in an appropriate spool directory and tells the remote user to pick it up with <code>uupick</code> . See the <code>uuto(1C)</code> man page for details.
<code>uupick</code>	Retrieves files in <code>/var/spool/uucppublic/receive</code> when files are transferred to a computer by using <code>uuto</code> . See the <code>uuto(1C)</code> man page.
<code>uux</code>	Creates the work, data, and execute files that are needed to execute commands on a remote machine. See the <code>uux(1C)</code> man page for details.
<code>uustat</code>	Displays the status of requested transfers (<code>uucp</code> , <code>uuto</code> , or <code>uux</code>). <code>uustat</code> also provides a means of controlling queued transfers. See the <code>uustat(1C)</code> man page for details.

UUCP Database Files

A major part of UUCP setup is the configuration of the files that compose the UUCP database. These files are in the `/etc/uucp` directory. You need to edit these files to set up UUCP or `asppp` on your machine. The files include the following:

<code>Config</code>	Contains a list of variable parameters. You can manually set these parameters to configure the network.
<code>Devconfig</code>	Used to configure network communications.
<code>Devices</code>	Used to configure network communications.

Dialcodes	Contains dial-code abbreviations that can be used in the phone number field of <i>Systems</i> file entries. Though not required, <i>Dialcodes</i> can be used by <i>asppp</i> as well as UUCP.
Dialers	Contains character strings that are required to negotiate with modems to establish connections with remote computers. <i>Dialers</i> is used by <i>asppp</i> as well as UUCP.
Grades	Defines job grades, and the permissions that are associated with each job grade, which users can specify to queue jobs to a remote computer.
Limits	Defines the maximum number of simultaneous <i>uucicos</i> , <i>uuxqts</i> , and <i>uuscheds</i> that are permitted on your machine.
Permissions	Defines the level of access that is granted to remote hosts that attempt to transfer files or execute commands on your machine.
Poll	Defines machines that are to be polled by your system and when they are polled.
Sysfiles	Assigns different or multiple files to be used by <i>uucico</i> and <i>cu</i> as <i>Systems</i> , <i>Devices</i> , and <i>Dialers</i> files.
Sysname	Enables you to define a unique UUCP name for a machine, in addition to its TCP/IP host name.
Systems	Contains information that is needed by the <i>uucico</i> daemon, <i>cu</i> , and <i>asppp</i> to establish a link to a remote computer. This information includes the name of the remote host, the name of the connecting device associated with the remote host, time when the host can be reached, telephone number, login ID, and password.

Several other files can be considered part of the supporting database but are not directly involved in establishing a link and transferring files.

Configuring UUCP Database Files

The UUCP database consists of the files that are shown in “UUCP Database Files” on page 574. However, basic UUCP configuration involves only the following critical files:

- */etc/uucp/Systems*
- */etc/uucp/Devices*
- */etc/uucp/Dialers*

Because *asppp* uses some of the UUCP databases, you should understand at minimum these critical database files if you plan to configure *asppp*. After these databases are configured, UUCP administration is fairly straightforward. Typically,

you edit the `Systems` file first, then edit the `Devices` file. You can usually use the default `/etc/uucp/Dialers` file, unless you plan to add dialers that aren't in the default file. In addition, you might also want to use the following files for basic UUCP and `asppp` configuration:

- `/etc/uucp/Sysfiles`
- `/etc/uucp/Dialcodes`
- `/etc/uucp/Sysname`

Because these files work closely with one another, you should understand all their contents before you change any one of them. A change to an entry in one file might require a change to a related entry in another file. The remaining files that are listed in "UUCP Database Files" on page 574 are not as critically intertwined.

Note – `asppp` uses only the files that are described in this section. `asppp` does not use the other UUCP database files.

Administering UUCP

This chapter explains how to start UUCP operations after you have modified the database file relevant to your machines. The chapter contains procedures and troubleshooting information for setting up and maintaining UUCP on machines that run the Solaris environment, such as the following:

- “UUCP Administration Task Map” on page 577
- “Adding UUCP Logins” on page 578
- “Starting UUCP” on page 579
- “Running UUCP Over TCP/IP” on page 581
- “UUCP Security and Maintenance” on page 582
- “Troubleshooting UUCP” on page 583

UUCP Administration Task Map

The following table provides pointers to the procedures that are covered in this chapter, in addition to a short description of each procedure.

TABLE 39-1 Task Map: UUCP Administration

Task	Description	For Instructions
Allow remote machines to have access to your system	Edit the <code>/etc/passwd</code> file to add entries to identify the machines that are permitted to access your system.	“How to Add UUCP Logins” on page 578
Start UUCP	Use the supplied shell scripts to start UUCP.	“How to Start UUCP” on page 579

TABLE 39-1 Task Map: UUCP Administration (Continued)

Task	Description	For Instructions
Enable UUCP to work with TCP/IP	Edit <code>/etc/inetd.conf</code> and <code>/etc/uucp/Systems</code> files to activate UUCP for TCP/IP.	"How to Activate UUCP for TCP/IP" on page 581
Troubleshoot some common UUCP problems	Diagnostic steps to use to check for faulty modems or ACUs. Diagnostic steps to use for debugging transmissions.	"How to Check for Faulty Modems or ACUs" on page 583 "How to Debug Transmissions" on page 583

Adding UUCP Logins

For incoming UUCP (`uucico`) requests from remote machines to be handled properly, each machine has to have a login on your system.

▼ How to Add UUCP Logins

To allow a remote machine to access your system, you need to add an entry to the `/etc/passwd` file as follows:

- 1. Edit the `/etc/passwd` file and add the entry to identify the machine that is permitted to access your system.**

A typical entry that you might put into the `/etc/passwd` file for a remote machine that is permitted to access your system with a UUCP connection would be as follows:

```
Ugobi:*:5:5:gobi:/var/spool/uucppublic:/usr/lib/uucp/uucico
```

By convention, the login name of a remote machine is the machine name preceded by the uppercase letter U. Note that the name should not exceed eight characters, so that in some situations you might have to truncate or abbreviate it.

The previous entry shows that a login request by Ugobi is answered by `/usr/lib/uucp/uucico`. The home directory is `/var/spool/uucppublic`. The password is obtained from the `/etc/shadow` file. You must coordinate the password and the login name with the UUCP administrator of the remote machine. The remote administrator must then add an appropriate entry, with login name and unencrypted password, in the remote machine's `Systems` file.

- 2. Coordinate your machine name with the UUCP administrators on other systems.**

Similarly, you must coordinate your machine's name and password with the UUCP administrators of all machines that you want to reach through UUCP.

Starting UUCP

UUCP includes four shell scripts that poll remote machines, reschedule transmissions, and clean up old log files and unsuccessful transmissions. The scripts are as follows:

- `uudemon.poll`
- `uudemon.hour`
- `uudemon.admin`
- `uudemon.cleanup`

These shell scripts should execute regularly to ensure that UUCP runs smoothly. The `crontab` file to run the scripts is automatically created in `/usr/lib/uucp/uudemon.crontab` as part of the Solaris installation process, if you select the full installation. Otherwise, the file is created when you install the UUCP package.

You can also run the UUCP shell scripts manually. The following is the prototype `uudemon.crontab` file that you can tailor for a particular machine:

```
#
#ident    "@(#)uudemon.crontab    1.5    97/12/09 SMI"
#
# This crontab is provided as a sample. For systems
# running UUCP edit the time schedule to suit, uncomment
# the following lines, and use crontab(1) to activate the
# new schedule.
#
#48 8,12,16 * * * /usr/lib/uucp/uudemon.admin
#20 3 * * * /usr/lib/uucp/uudemon.cleanup
#0 * * * * /usr/lib/uucp/uudemon.poll
#11,41 * * * * /usr/lib/uucp/uudemon.hour
```

Note – By default, UUCP operations are disabled. To enable UUCP, edit the time schedule and uncomment the appropriate lines in the `uudemon.crontab` file.

▼ How to Start UUCP

To activate the `uudemon.crontab` file, do the following:

1. **Become superuser.**
2. **Edit the `/usr/lib/uucp/uudemon.crontab` file and change entries as required.**
3. **Activate the `uudemon.crontab` file by issuing the following command:**

```
crontab < /usr/lib/uucp/uudemon.crontab
```

uudemon.poll Shell Script

The default `uudemon.poll` shell script reads the `/etc/uucp/Poll` file once an hour. If any machines in the `Poll` file are scheduled to be polled, a work file (`C.synxxx`) is placed in the `/var/spool/uucp/nodename` directory, where *nodename* represents the UUCP node name of the machine.

The shell script is scheduled to run once an hour, before `uudemon.hour`, so that the work files are in place when `uudemon.hour` is called.

uudemon.hour Shell Script

The default `uudemon.hour` shell script does the following:

- Calls the `uusched` program to search the spool directories for work files (`C.`) that have not been processed. The script then schedules these files for transfer to a remote machine.
- Calls the `uuxqt` daemon to search the spool directories for execute files (`X.`) that have been transferred to your computer and were not processed at the time they were transferred.

By default, `uudemon.hour` runs twice an hour. You might want it to run more often if you expect high failure rates of calls to remote machines.

uudemon.admin Shell Script

The default `uudemon.admin` shell script does the following:

- Runs the `uustat` command with `p` and `q` options. The `q` reports on the status of work files (`C.`), data files (`D.`), and execute files (`X.`) that are queued. The `p` prints process information for networking processes that are listed in the lock files (`/var/spool/locks`).
- Sends resulting status information to the uucp administrative login by using `mail`.

uudemon.cleanup Shell Script

The default `uudemon.cleanup` shell script does the following:

- Collects log files for individual machines from the `/var/uucp/.Log` directory, merges them, and places them in the `/var/uucp/.Old` directory with other old log information
- Removes work files (`C.`) seven days old or older, data files (`D.`) seven days old or older, and execute files (`X.`) two days old or older from the spool files

- Returns mail that cannot be delivered to the sender
- Mails a summary of the status information that was gathered during the current day to the UUCP administrative login (uucp)

Running UUCP Over TCP/IP

To run UUCP on a TCP/IP network, you need to make a few modifications, as described in this section.

▼ How to Activate UUCP for TCP/IP

1. **Edit the `/etc/inetd.conf` file and ensure that the following entry is not preceded by a comment mark (#):**

```
uucp stream tcp nowait root /usr/sbin/in.uucpd in.uucpd
```

2. **Edit the `/etc/uucp/Systems` file to ensure that the entries have the following fields :**

System-Name Time TCP Port networkname Standard-Login-Chat

A typical entry would resemble the following:

```
rochester Any TCP - ur-seneca login: Umachine password: xxx
```

Notice that the *networkname* field permits you to specify explicitly the TCP/IP host name. This is important for some sites. In the previous example, the site has the UUCP node name *rochester*, which is different from its TCP/IP host name *ur-seneca*. Moreover, a completely different machine could easily run UUCP and have the TCP/IP host name of *rochester*.

The Port field in the Systems file should have the entry `-`. This syntax is equivalent to listing the entry as `uucp`. In almost every situation, the *networkname* is the same as the system name, and the Port field is `-`, which says to use the standard uucp port from the services database. The `in.uucpd` daemon expects the remote machine to send its login and password for authentication, and `in.uucpd` prompts for them, much as `getty` and `login` do.

3. **Edit the `/etc/inet/services` file to set up a port for UUCP:**

```
uucp      540/tcp      uucpd      # uucp daemon
```

You should not have to change the entry. However, if your machine runs NIS or NIS+ as its name service, you should change the `/etc/nsswitch.conf` entry for `/etc/services` to check files first, then check `nis` or `nisplus`.

UUCP Security and Maintenance

After you have set up UUCP, maintenance is straightforward. This section explains ongoing UUCP tasks that relate to security, maintenance, and troubleshooting.

Setting Up UUCP Security

The default `/etc/uucp/Permissions` file provides the maximum amount of security for your UUCP links. The default `Permissions` file contains no entries.

You can set additional parameters for each remote machine to define the following:

- Ways the remote machine can receive files from your machine
- Directories for which the remote machine has read and write permission
- Commands the remote machine can use for remote execution

A typical `Permissions` entry follows:

```
MACHINE=datsum LOGNAME=Udatsum VALIDATE=datsum  
COMMANDS=rmail REQUEST=yes SENDFILES=yes
```

This entry allows files to be sent and received to and from the “normal” UUCP directories, not from anywhere in the system. The entry also causes the UUCP user name to be validated at login time.

Regular UUCP Maintenance

UUCP does not require much maintenance. Except for ensuring that the `crontab` file is in place, as described in the section “How to Start UUCP” on page 579, your concern should be the growth of mail files and the public directory.

Email for UUCP

All email messages that are generated by the UUCP programs and scripts are sent to the user ID `uucp`. If you do not log in frequently as that user, you might not realize that mail is accumulating (and consuming disk space). To solve this problem, create an alias in `/etc/mail/aliases` and redirect that email either to `root` or to yourself and others responsible for maintaining UUCP. Remember to run the `newaliases` command after modifying the `aliases` file.

UUCP Public Directory

The directory `/var/spool/uucppublic` is the one place in every system to which UUCP by default is able to copy files. Every user has permission to change to `/var/spool/uucppublic` and read and write files in it. However, its sticky bit is set, so its mode is `01777`. As a result, users cannot remove files that have been copied to it and that belong to `uucp`. Only you, as UUCP administrator logged in as `root` or `uucp`, can remove files from this directory. To prevent the uncontrolled accumulation of files in this directory, you should ensure that you remove files from it periodically.

If this maintenance is inconvenient for users, encourage them to use `uuto` and `uupick` rather than removing the sticky bit, which is set for security reasons. See the `uuto(1C)` man page for instructions for using `uuto` and `uupick`. You can also restrict the mode of the directory to only one group of people. If you do not want to run the risk of someone filling your disk, you can even deny UUCP access to it.

Troubleshooting UUCP

These procedures describe how to solve common UUCP problems.

▼ How to Check for Faulty Modems or ACUs

You can check if the modems or other ACUs are not working properly in several ways.

1. **Obtain counts and reasons for contact failure by running the following command:**

```
# uustat -q
```

2. **Call over a particular line and print debugging information on the attempt.**

The line must be defined as `direct` in the `/etc/uucp/Devices` file. (You must add a telephone number to the end of the command line if the line is connected to an autodialer or the device must be set up as `direct`.) Type:

```
# cu -d -lline  
line is /dev/cua/a.
```

▼ How to Debug Transmissions

If you cannot contact a particular machine, you can check out communications to that machine with `Uutry` and `uucp`.

1. **To try to make contact:**

```
# /usr/lib/uucp/Uutry -r machine
```

Replace *machine* with the host name of the machine you are having problems contacting. This command does the following:

- Starts the transfer daemon (*uucico*) with debugging. You can get more debugging information if you are *root*.
- Directs the debugging output to */tmp/machine*.
- Prints the debugging output to your terminal by issuing the following command:

```
# tail -f
```

Press Control-c to end output. You can copy the output from */tmp/machine* if you want to save it.

2. If *Uutry* doesn't isolate the problem, try to queue a job:

```
# uucp -r file machine\1/dir/file
```

Replace *file* by the file you want to transfer, *machine* by the machine you want to copy to, and */dir/file* where the file will be placed on the other machine. The *r* option queues a job but does not start the transfer.

3. Issue the following command:

```
# Uutry
```

If you still cannot solve the problem, you might need to call your local support representative. Save the debugging output. It can help diagnose the problem.

You might also decrease or increase the level of debugging that is provided by *Uutry* through the *-x n* option. *n* indicates the debug level. The default debug level for *Uutry* is 5.

Debug level 3 provides basic information about when and how the connection is established, but not much information about the transmission itself. Debug level 9, however, provides exhaustive information about the transmission process. Be aware that debugging occurs at both ends of the transmission. If you intend to use a level higher than 5 on a moderately large text, contact the administrator of the other site and agree on a time for doing so.

Checking the UUCP */etc/uucp/Systems* File

Verify that you have up-to-date information in your *Systems* file if you are having trouble contacting a particular machine. Some information that might be out of date for a machine are the following:

- Phone number
- Login ID
- Password

Checking UUCP Error Messages

UUCP has two types of error messages: `ASSERT` and `STATUS`.

- When a process is aborted, `ASSERT` error messages are recorded in `/var/uucp/.Admin/errors`. These messages include the file name, `sccsid`, line number, and text. These messages usually result from system problems.
- `STATUS` error messages are stored in the `/var/uucp/.Status` directory. The directory contains a separate file for each remote machine your computer attempts to communicate with. These files contain status information on the attempted communication and whether it was successful.

Checking Basic Information

Several commands are available for checking basic networking information:

- Use the `uname` command to list those machines your machine can contact.
- Use the `uulog` command to display the contents of the log directories for particular hosts.
- Use the `uucheck -v` command to check for the presence of files and directories that are needed by `uucp`. This command also checks the `Permissions` file and outputs information on the permissions you have set up.

Composed September 5, 2002

UUCP Reference

This chapter provides reference information for working with UUCP. The following topics are covered:

- “UUCP /etc/uucp/Systems File” on page 587
- “UUCP /etc/uucp/Devices File” on page 593
- “UUCP /etc/uucp/Dialers File” on page 599
- “Other Basic UUCP Configuration Files” on page 603
- “UUCP /etc/uucp/Permissions File” on page 605
- “UUCP /etc/uucp/Poll File” on page 613
- “UUCP /etc/uucp/Config File” on page 613
- “UUCP /etc/uucp/Grades File” on page 614
- “Other UUCP Configuration Files” on page 616
- “UUCP Administrative Files” on page 618
- “UUCP Error Messages” on page 619

UUCP /etc/uucp/Systems File

The `/etc/uucp/Systems` file contains the information that is needed by the `uucico` daemon to establish a communication link to a remote computer.

`/etc/uucp/Systems` is the first file you need to edit to configure UUCP.

Each entry in the `Systems` file represents a remote computer with which your host communicates. A particular host can have more than one entry. The additional entries represent alternative communication paths that are tried in sequential order. In addition, by default UUCP prevents any computer that does not appear in `/etc/uucp/Systems` from logging in to your host.

Using the `Sysfiles` file, you can define several files to be used as `Systems` files. See “UUCP /etc/uucp/Sysfiles File” on page 604 for a description of `Sysfiles`.

Each entry in the `Systems` file has the following format:

<i>System-Name</i>	<i>Time</i>	<i>Type</i>	<i>Speed</i>	<i>Phone</i>	<i>Chat Script</i>

The following example shows the fields of the `Systems` file.

EXAMPLE 40-1 Fields in `/etc/uucp/Systems`

System-Name	Time	Type	Speed	Phone	Chat Script
Arabian	Any	ACUEC	38400	111222	Login: Puucp ssword:beledi

UUCP System-Name Field

This field contains the node name of the remote computer. On TCP/IP networks, this name can be the machine's host name or a name that is created specifically for UUCP communications through the `/etc/uucp/Sysname` file. See "UUCP `/etc/uucp/Systems` File" on page 587. In Example 40-1, the `System-Name` field contains an entry for remote host `arabian`.

UUCP Time Field

This field specifies the day of week and time of day when the remote computer can be called. The format of the `Time` field follows:

`daytime[;retry]`

The *day* portion can be a list that contains some of the following entries.

TABLE 40-1 Day Field

Su Mo Tu We Th Fr Sa	For individual days.
Wk	For any weekday.
Any	For any day.
Never	Your host never initiates a call to the remote computer. The call must be initiated by the remote computer. Your host is then operating in <i>passive mode</i> .

Example 40-1 shows `Any` in the `Time` field, which indicates that host `arabian` can be called at any time.

The *time* portion should be a range of times that are specified in 24-hour notation, for example, `0800-1230` for 8:30 a.m. to 12:30 p.m. If no *time* portion is specified, any time of day is assumed to be allowed for the call.

A time range that spans 0000 is permitted. For example, 0800-0600 means all times are allowed other than times between 6 a.m. and 8 a.m.

UUCP Retry Subfield

The *Retry* subfield enables you to specify the minimum time (in minutes) before a retry, following a failed attempt. The default wait is 60 minutes. The subfield separator is a semicolon (;). For example, Any; 9 is interpreted as call any time, but wait at least 9 minutes before retrying after a failure occurs.

If you do not specify a *retry* entry, an exponential back-off algorithm is used. This means that UUCP starts with a default wait time that grows larger as the number of failed attempts increases. For example, suppose the initial retry time is 5 minutes. If no response occurs, the next retry is 10 minutes later. The next retry is 20 minutes later, and so on until the maximum retry time of 23 hours is reached. If *retry* is specified, that is always the retry time. Otherwise, the back-off algorithm is used.

UUCP Type Field

This field contains the device type that should be used to establish the communication link to the remote computer. The keyword that is used in this field is matched against the first field of *Devices* file entries.

EXAMPLE 40-2 Type Field and /etc/uucp/Devices File

File Name System-Name Time Type Speed Phone Chat Script

```
Systems    arabian      Any    ACUEC, g 38400 1112222  ogin: Puucp ssword:beledi
```

You can define the protocol that is used to contact the system by adding it on to the Type field. The previous example shows how to attach the protocol *g* to the device type *ACUEC*. For information on protocols, see “UUCP Protocol Definitions in the *Devices* File” on page 598.

UUCP Speed Field

This field (also known as the Class field) specifies the transfer speed of the device that is used in establishing the communication link. The UUCP speed field can contain a letter and speed (for example, C1200, D1200) to differentiate between classes of dialers (refer to “UUCP Class Field” on page 595).

Some devices can be used at any speed, so the keyword *Any* can be used. This field must match the Class field in the associated *Devices* file entry.

EXAMPLE 40-3 Speed Field and /etc/uucp/Devices File

File Name System-Name Time Type Speed Phone Chat Script

EXAMPLE 40-3 Speed Field and /etc/uucp/Devices File (Continued)

```
Systems    eagle          Any    ACU, g    D1200    NY3251    ogin: nuucp  ssword: Oakgrass
```

If information is not required for this field, use a dash (-) as a placeholder for the field.

UUCP Phone Field

This field enables you to specify the telephone number (token) of the remote computer for automatic dialers (port selectors). The telephone number consists of an optional alphabetic abbreviation and a numeric part. If an abbreviation is used, it must be one that is listed in the `Dialcodes` file.

EXAMPLE 40-4 Phone Field Correspondence

File Name	System-Name	Time	Type	Speed	Phone	Chat Script

Systems	nubian		Any	ACU	2400	NY5551212 ogin: Puucp ssword:Passuan

In the `System-Name` string, an equals sign (=) tells the ACU to wait for a secondary dial tone before dialing the remaining digits. A dash (-) in the string instructs the ACU to pause four seconds before dialing the next digit.

If your computer is connected to a port selector, you can access other computers that are connected to that selector. The `Systems` file entries for these remote machines should not have a telephone number in the `Phone` field. Instead, this field should contain the token to be passed on to the switch. In this way, the port selector knows the remote machine with which your host wants to communicate, usually just the system name. The associated `Devices` file entry should have a \D at the end of the entry to ensure that this field is not translated by using the `Dialcodes` file.

UUCP Chat Script Field

This field (also called the Login field) contains a string of characters that is called a *chat-script*. The chat-script contains the characters the local and remote machines must pass to each other in their initial conversation. Chat-scripts have the following format:

expect send [expect send] ...

expect represents the string that the local host expects to receive from the remote host to initiate conversation. *send* is the string the local host sends after it receives the *expect* string from the remote host. A chat-script can have more than one expect-send sequence.

A basic chat-script might contain the following:

- Login prompt that the local host expects to receive from the remote machine
- Login name that the local host sends to the remote machine in order to log in
- Password prompt that the local host expects to receive from the remote machine
- Password that the local host sends to the remote machine

The *expect* field can be composed of subfields of the following form:

expect[-send-expect]...

send is sent if the prior *expect* is not successfully read, and the *-expect* that follows the *send* is the next expected string.

For example, with strings `login--login`, the UUCP on the local host expects `login`. If UUCP receives `login` from the remote machine, it goes to the next field. If it does not receive `login`, it sends a carriage return, then looks for `login` again. If the local computer initially does not expect any characters, use the characters `" "` (NULL string) in the *expect* field. All *send* fields are sent with a carriage return appended unless the *send* string is terminated with a `\c`.

Here is an example of a Systems file entry that uses an *expect-send* string:

```
System-Name Time Type Speed Phone Chat Script
sonora Any ACUEC 9600 2223333 " " \r \r ogin:-BREAK-ogin: Puucpx ssword: xyzzy
```

This example tells UUCP on the local host to send two carriage returns and wait for `ogin:` (for `Login:`). If `ogin:` is not received, send a `BREAK`. When you do receive `ogin:` send the login name `Puucpx`. When you receive `ssword:` (for `Password:`), send the password `xyzzy`.

The following table lists some useful escape characters.

TABLE 40-2 Escape Characters Used in Systems File Chat Script

Escape Character	Meaning
<code>\b</code>	Sends or expects a backspace character.
<code>\c</code>	If at the end of a string, suppresses the carriage return that is normally sent. Ignored otherwise.
<code>\d</code>	Delays 1–3 seconds before sending more characters.
<code>\E</code>	Starts echo checking. From this point on, whenever a character is transmitted, UUCP waits for the character to be received before continuing its checks.
<code>\e</code>	Echoes check-off.
<code>\H</code>	Ignores one hangup. Use this option for dialback modems.
<code>\K</code>	Sends a <code>BREAK</code> character.
<code>\M</code>	Turns on <code>CLOCAL</code> flag.

TABLE 40-2 Escape Characters Used in Systems File Chat Script *(Continued)*

Escape Character	Meaning
\m	Turns off CLOCAL flag.
\n	Sends or expects a newline character.
\N	Sends a NULL character (ASCII NUL).
\p	Pauses for approximately 1/4 to 1/2 second.
\r	Sends or expects a carriage return.
\s	Sends or expects a space character.
\t	Sends or expects a tab character.
EOT	Sends an EOT, followed by newline twice.
BREAK	Sends a break character.
\ddd	Sends or expects the character that is represented by the octal digits (<i>ddd</i>).

Enabling Dialback Through the Chat Script

Some companies set up dial-in servers to handle calls from remote computers. For example, your company might have a dial-in server with a dialback modem that employees can call from their home computers. After the dial-in server identifies the remote machine, it disconnects the link to the remote machine and then calls back the remote machine. The communications link is then reestablished.

You can facilitate dialback by using the \H option in the Systems file chat-script at the place where dialback should occur. Include the \H as part of an expect string at the place where the dial-in server is expected to hang up.

For example, suppose the chat-script that calls a dial-in server contains the following string:

```
INITIATED\Hogin:
```

The UUCP dialing facility on the local machine expects to receive the characters INITIATED from the dial-in server. After the INITIATED characters have been matched, the dialing facility flushes any subsequent characters it receives until the dial-in server hangs up. The local dialing facility then waits until it receives the next part of the expect string, the characters ogin:, from the dial-in server. When it receives the ogin:, the dialing facility then continues through the chat-script.

You need not to have a string of characters directly preceding or following the \H, as shown in the previous sample string.

UUCP Hardware Flow Control

You can also use the pseudo-send `STTY=value` string to set modem characteristics. For instance, `STTY=crtscts` enables hardware flow control. `STTY` accepts all `stty` modes. See the `stty(1)` and `termio(7I)` man pages for complete details.

The following example would enable hardware flow control in a `Systems` file entry:

```
System-Name Time Type Speed Phone Chat Script
unix Any ACU 2400 12015551212 "" \r login:-\r-login:-\r-login:
nuucp password: xxx "" \ STTY=crtscts
```

This pseudo-send string can also be used in entries in the `Dialers` file.

UUCP Setting Parity

In some situations, you have to reset the parity because the system that you are calling checks port parity and drops the line if it is wrong. The expect-send couplet `" P_ZERO` sets the high-order bit (parity bit) to 0. For example:

```
System-Name Time Type Speed Phone Chat Script
unix Any ACU 2400 12015551212 "" P_ZERO "" \r login:-\r-login:-\r-login:
nuucp password: xxx
```

In the same manner, `P_EVEN` sets parity to even (the default), `P_ODD` sets odd parity, and `P_ONE` sets the parity bit to 1.

The parity couplet can be inserted anywhere in the chat-script. The couplet applies to all information in the chat-script that follows the `" P_ZERO`. The couplet can also be used in entries in the `Dialers` file.

UUCP /etc/uucp/Devices File

The `/etc/uucp/Devices` file contains information for all the devices that can be used to establish a link to a remote computer. These devices include ACUs—which includes modern, high-speed modems—direct links, and network connections.

Here is an entry in `/etc/uucp/Devices` for a US Robotics V.32bis modem that is attached to port A and running at 38,400 bps.

```
Type Line Line2 Class Dialer-Token-Pairs
ACUEC cua/a - 38400 usrv32bis-ec
```

Each field is described in the next section.

UUCP Type Field

This field describes the type of link that the device establishes. The UUCP Type field can contain one of the keywords that is described in the sections that follow.

Direct Keyword

The `Direct` keyword appears mainly in entries for `cu` connections. This keyword indicates that the link is a direct link to another computer or a port selector. Create a separate entry for each line that you want to reference through the `-l` option of `cu`.

ACU Keyword

The `ACU` keyword indicates that the link to a remote computer (whether through `cu`, `UUCP`, `asppp`, or Solaris PPP 4.0) is made through a modem. This modem can be connected either directly to your computer or indirectly through a port selector.

Port Selector

This is a variable that is replaced in the Type field by the name of a port selector. Port selectors are devices that are attached to a network that prompts for the name of a calling modem, then grant access. The file `/etc/uucp/Dialers` contains caller scripts only for the `micom` and `develcon` port selectors. You can add your own port selector entries to the `Dialers` file. See “UUCP `/etc/uucp/Dialers` File” on page 599 for more information.

Sys-Name

This variable is replaced by the name of a machine in the Type field, indicating that the link is a direct link to this particular computer. This naming scheme is used to associate the line in this `Devices` entry with an entry in `/etc/uucp/Systems` for the computer *Sys-Name*.

Type Field and `/etc/uucp/Systems` File

Example 40–5 shows a comparison between the fields in `/etc/uucp/Devices` and fields in `/etc/uucp/Systems`. The titles of each column apply only to fields in the `Devices` file.

The keyword that is used in the Type field of the `Devices` file is matched against the third field of the `Systems` file entries. In the `Devices` file, the Type field has the entry `ACUEC`, indicating an automatic call unit, in this instance a V.32bis modem. This value is matched against the third field in the `Systems` file, which also contains the entry `ACUEC`. See “UUCP `/etc/uucp/Systems` File” on page 587 for more information.

Composed September 5, 2002

EXAMPLE 40-5 Type Field and /etc/uucp/Systems File Equivalent

File Name Type Line Line2 Class Dialer-Token-Pairs

Devices ACUEC cua/a - 38400 usrv32bis-ec

System nubian Any ACUEC 38400 9998888 "" \d\d\r\n\c-ogin-\r\n\c-ogin.....

UUCP Line Field

This field contains the device name of the line (port) that is associated with the `Devices` entry. If the modem that is associated with a particular entry were attached to the `/dev/cua/a` device (serial port A), the name that is entered in this field would be `cua/a`. An optional modem control flag, `M`, can be used in the Line field to indicate that the device should be opened without waiting for a carrier. For example:

`cua/a,M`

UUCP Line2 Field

This field is a placeholder. Always use a dash (-) here. 801-type dialers, which are not supported in the Solaris environment, use the Line2 field. Non-801 dialers do not normally use this configuration, but still require a hyphen in this field.

UUCP Class Field

The Class field contains the speed of the device, if the keyword `ACU` or `Direct` is used in the Type field. However, the Class field can contain a letter and a speed (for example, `C1200`, `D1200`) to differentiate between classes of dialers (Centrex or Dimension PBX).

This differentiation is necessary because many larger offices can have more than one type of telephone network: one network might be dedicated to serving only internal office communications while another handles the external communications. In such a situation, you must distinguish which line(s) should be used for internal communications and which should be used for external communications.

The keyword that is used in the Class field of the `Devices` file is matched against the Speed field of `Systems` file.

EXAMPLE 40-6 UUCP Class Field

File Name Type Line Line2 Class Dialer-Token-Pairs

Devices ACU cua/a - D2400 hayes

Some devices can be used at any speed, so the keyword *Any* can be used in the *Class* field. If *Any* is used, the line matches any speed that is requested in the *Speed* field of the *Systems* file. If this field is *Any* and the *Systems* file *Speed* field is *Any*, the speed defaults to 2400 bps.

UUCP Dialer-Token-Pairs Field

The Dialer-Token-Pairs (DTP) field contains the name of a dialer and the token to pass it. The DTP field has this syntax:

dialer token [dialer token]

The *dialer* portion can be the name of a modem, a port monitor, or it can be *direct* or *uudirect* for a direct-link device. You can have any number of dialer-token pairs. If the *dialer* portion is not present, it is taken from a related entry in the *Systems* file. The *token* portion can be supplied immediately after the dialer portion.

The last dialer-token pair might not be present, depending on the associated dialer. In most situations, the last pair contains only a *dialer* portion. The *token* portion is retrieved from the *Phone* field of the associated *Systems* file entry.

A valid entry in the *dialer* portion can be defined in the *Dialers* file or can be one of several special dialer types. These special dialer types are compiled into the software and are therefore available without having entries in the *Dialers* file. The following table shows the special dialer types.

TABLE 40-3 Dialer-Token Pairs

TCP	TCP/IP network
TLI	Transport Level Interface Network (without STREAMS)
TLIS	Transport Level Interface Network (with STREAMS)

See “UUCP Protocol Definitions in the *Devices* File” on page 598 for more information.

Structure of the Dialer-Token-Pairs Field

The DTP field can be structured four different ways, depending on the device that is associated with the entry:

- **Directly connected modem**

If a modem is connected directly to a port on your computer, the DTP field of the associated *Devices* file entry has only one pair. This pair would normally be the name of the modem. This name is used to match the particular *Devices* file entry with an entry in the *Dialers* file. Therefore, the *Dialer* field must match the first

field of a Dialers file entry.

EXAMPLE 40-7 Dialers Field for Direct Connect Modem

```
Dialers   hayes =,-,  ""           \\dA\pTE1V1X1Q0S2=255S12=255\r\c
                                         \EATDT\T\r\c CONNECT
```

Notice that only the dialer portion (hayes) is present in the DTP field of the *Devices* file entry. This means that the *token* to be passed on to the dialer (in this instance, the phone number) is taken from the *Phone* field of a *Systems* file entry. (\T is implied, as described in Example 40-9.)

- **Direct link** – For a direct link to a particular computer, the DTP field of the associated entry would contain the keyword *direct*. This is true for both types of direct-link entries, *Direct* and *Sys-Name* (refer to “UUCP Type Field” on page 594).
- **Computers on the same port selector** – If a computer with which you intend to communicate is on the same port selector switch as your computer, your computer must first access the switch. The switch then makes the connection to the other computer. This type of entry has only one pair. The *dialer* portion is used to match a *Dialers* file entry.

EXAMPLE 40-8 UUCP Dialers Field for Computers on Same Port Selector

```
Dialers   develcon ,""  ""           \pr\ps\c est:\007 \E\D\e \007
```

As shown, the *token* portion is left blank. This designation indicates that it is retrieved from the *Systems* file. The *Systems* file entry for this computer contains the token in the *Phone* field, which is normally reserved for the phone number of the computer. Refer to “UUCP /etc/uucp/Systems File” on page 587 for details. This type of DTP contains an escape character (\D), which ensures that the contents of the *Phone* field are not interpreted as a valid entry in the *Dialcodes* file.

- **Modems that are connected to port selector** – If a high-speed modem is connected to a port selector, your computer must first access the port selector switch. The switch makes the connection to the modem. This type of entry requires two dialer-token-pairs. The *dialer* portion of each pair (fifth and seventh fields of entry) is used to match entries in the *Dialers* file, as shown below.

EXAMPLE 40-9 UUCP Dialers Field for Modems Connected to Port Selector

```
Dialers   develcon ""  ""  \pr\ps\c est:\007  \E\D\e  \007
Dialers   ventel  =&-%  t""  \r\p\r\c $  <K\T%\r>\c  ONLINE!
```

In the first pair, *develcon* is the dialer and *vent* is the token that is passed to the *Develcon* switch to tell it which device (such as *Ventel* modem) to connect to your computer. This token is unique for each port selector, as each switch can be set up differently. After the *Ventel* modem has been connected, the second pair is accessed. *Ventel* is the dialer and the token is retrieved from the *Systems* file.

Two escape characters can appear in a DTP field:

- \T – Indicates that the Phone (*token*) field should be translated by using the `/etc/uucp/Dialcodes` file. This escape character is normally placed in the `/etc/uucp/Dialers` file for each caller script that is associated with a modem (Hayes, US Robotics, and so on). Therefore, the translation does not occur until the caller script is accessed.
- \D – Indicates that the Phone (*token*) field should not be translated by using the `/etc/uucp/Dialcodes` file. If no escape character is specified at the end of a Devices entry, the \D is assumed (default). A \D is also used in the `/etc/uucp/Dialers` file with entries that are associated with network switches (develcon and micom).

UUCP Protocol Definitions in the Devices File

You can define the protocol to use with each device in `/etc/uucp/Devices`. This specification is usually unnecessary because you can use the default or define the protocol with the particular system you are calling. Refer to “UUCP `/etc/uucp/Systems` File” on page 587 for details. If you do specify the protocol, you must use the following form:

Type,Protocol [parameters]

For example, you can use `TCP,te` to specify the TCP/IP protocol.

The following table shows the available protocols for the `Devices` file.

TABLE 40–4 Protocols Used in `/etc/uucp/Devices`

Protocol	Description
t	This protocol is commonly used for transmissions over TCP/IP and other reliable connections. t assumes error-free transmissions.
g	This protocol is UUCP’s native protocol. It is slow, reliable, and good for transmission over noisy telephone lines.
e	This protocol assumes transmission over error-free channels that are message oriented (as opposed to byte-stream oriented, such as TCP/IP).
f	This protocol is used for transmission over X.25 connections. f relies on flow control of the data stream, and is meant for working over links that can (almost) be guaranteed to be error-free, specifically X.25/PAD links. A checksum is enacted over a whole file only. If a transport fails, the receiver can request retransmission(s).

Here is an example that shows a protocol designation for a device entry:

```
TCP,te - - Any TCP -
```

This example indicates that, for device TCP, you should try to use the `t` protocol. If the other end refuses, use the `e` protocol.

Neither `e` nor `t` is appropriate for use over modems. Even if the modem assures error-free transmission, data can still be dropped between the modem and the CPU.

UUCP /etc/uucp/Dialers File

The `/etc/uucp/Dialers` file contains dialing instructions for many commonly used modems. You probably do not need to change or add entries to this file unless you plan to use a nonstandard modem or plan to customize your UUCP environment. Nevertheless, you should understand what is in the file and how it relates to the `Systems` and `Devices` file.

The text specifies the initial conversation that must occur on a line before it can be made available for transferring data. This conversation, often referred to as a chat-script, is usually a sequence of ASCII strings that is transmitted and expected, and it is often used to dial a phone number.

As shown in the examples in “UUCP /etc/uucp/Devices File” on page 593, the fifth field in a `Devices` file entry is an index into the `Dialers` file or a special dialer type (TCP, TLI, or TLIS). The `uucico` daemon attempts to match the fifth field in the `Devices` file with the first field of each `Dialers` file entry. In addition, each odd-numbered `Devices` field, starting with the seventh position, is used as an index into the `Dialers` file. If the match succeeds, the `Dialers` entry is interpreted to perform the dialer conversation.

Each entry in the `Dialers` file has the following format:

<i>dialer</i>	<i>substitutions</i>	<i>expect-send</i>

The following example shows the entry for a US Robotics V.32bis modem.

EXAMPLE 40-10 /etc/uucp/Dialers File Entry

Dialer	Substitution	Expect-Send
usrv32bis-e	=, -, "	dA\pT&FE1V1X1Q0S2=255S12=255&A1&H1&M5&B2&W\r\c OK\r \EATDT\T\r\c CONNECT\s14400/ARQ STTY=crtscs

The `Dialer` field matches the fifth and additional odd-numbered fields in the `Devices` file. The `Substitutions` field is a translate string: the first of each pair of characters is mapped to the second character in the pair. This mapping is usually used to translate `=` and `-` into whatever the dialer requires for “wait for dial tone” and “pause.”

The remaining expect-send fields are character strings.

The following example shows some sample entries in the `Dialers` file, as distributed when you install UUCP as part of the Solaris installation program.

EXAMPLE 40-11 Excerpts From `/etc/uucp/Dialers`

```
penril      =W-P "" \d > Q\c : \d- > s\p9\c )-W\p\r\ds\p9\c-) y\c : \E\TP > 9\c OK

ventel      =&-% "" \r\p\r\c $ <K\T%\r>\c ONLINE!

vadic       =K-K "" \005\p *- \005\p-* \005\p-* D\p BER? \E\T\e \r\c LINE

develcon    "" "" \pr\ps\c est:\007

\E\D\e \n\007 micom "" "" \s\c NAME? \D\r\c GO

hayes       =,-, "" \dA\pTE1V1X1Q0S2=255S12=255\r\c OK\r \EATDT\T\r\c CONNECT

# Telebit TrailBlazer
tb1200      =W-, "" \dA\pA\pA\pTE1V1X1Q0S2=255S12=255S50=2\r\c OK\r
\EATDT\T\r\c CONNECT\s1200
tb2400      =W-, "" \dA\pA\pA\pTE1V1X1Q0S2=255S12=255S50=3\r\c OK\r
\EATDT\T\r\c CONNECT\s2400
tbfast      =W-, "" \dA\pA\pA\pTE1V1X1Q0S2=255S12=255S50=255\r\c OK\r
\EATDT\T\r\c CONNECT\sFAST

# USrobotics, Codes, and DSI modems

dsi-ec      =,-, "" \dA\pTE1V1X5Q0S2=255S12=255*E1*F3*M1*S1\r\c OK\r \EATDT\T\r\c
CONNECT\sEC STTY=crtscts,crtsexoff

dsi-nec     =,-, "" \dA\pTE1V1X5Q0S2=255S12=255*E0*F3*M1*S1\r\c OK\r \EATDT\T\r\c CONNECT
STTY=crtscts,crtsexoff

usrv32bis-ec =,-, "" \dA\pT&FE1V1X1Q0S2=255S12=255&A1&H1&M5&B2&W\r\c OK\r \EATDT\T\r\c
CONNECT\s14400/ARQ STTY=crtscts,crtsexoff

usrv32-nec  =,-, "" \dA\pT&FE1V1X1Q0S2=255S12=255&A0&H1&M0&B0&W\r\c OK\r \EATDT\T\r\c
CONNECT STTY=crtscts,crtsexoff

codex-fast  =,-, "" \dA\pT&C1&D2*MF0*AA1&R1&S1*DE15*FL3S2=255S7=40S10=40*TT5&W\r\c OK\r
\EATDT\T\r\c CONNECT\s38400 STTY=crtscts,crtsexoff

tb9600-ec   =W-, "" \dA\pA\pA\pTE1V1X1Q0S2=255S12=255S50=6\r\c OK\r
\EATDT\T\r\cCONNECT\s9600 STTY=crtscts,crtsexoff

tb9600-nec  =W-, "" \dA\pA\pA\pTE1V1X1Q0S2=255S12=255S50=6S180=0\r\c OK\r \EATDT\T\r\c
CONNECT\s9600 STTY=crtscts,crtsexoff
```

The following table lists escape characters that are commonly used in the send strings in the `Dialers` file.

TABLE 40-5 Backslash Characters for `/etc/uucp/Dialers`

Character	Description
<code>\b</code>	Sends or expects a backspace character.
<code>\c</code>	No newline or carriage return.
<code>\d</code>	Delays (approximately 2 seconds).
<code>\D</code>	Phone number or token without <code>Dialcodes</code> translation.
<code>\e</code>	Disables echo checking.
<code>\E</code>	Enables echo checking (for slow devices).
<code>\K</code>	Inserts a Break character.
<code>\n</code>	Sends newline.
<code>\nnn</code>	Sends octal number. Additional escape characters that can be used are listed in the section “UUCP <code>/etc/uucp/Systems File</code> ” on page 587.
<code>\N</code>	Sends or expects a NULL character (ASCII NUL).
<code>\p</code>	Pauses (approximately 12–14 seconds).
<code>\r</code>	Returns.
<code>\s</code>	Sends or expects a space character.
<code>\T</code>	Phone number or token with <code>Dialcodes</code> translation.

Here is a `penril` entry in the `Dialers` file:

```
penril =W-P "" \d > Q\c : \d- > s\p9\c )-W\p\r\ds\p9\c-) y\c : \E\TP > 9\c OK
```

First, the substitution mechanism for the phone number argument is established so that `any =` is replaced with a `W` (wait for dial tone) and any `-` with a `P` (pause).

The handshake that is given by the remainder of the line works as listed:

- `""` – Waits for nothing (that is, proceed to the next step).
- `\d` – Delays 2 seconds, then sends a carriage return.
- `>` – Waits for a `>`.
- `Q\c` – Sends a `Q` without a carriage return.
- `:` – Expects a `:`.
- `\d-` – Delays 2 seconds, sends a `-` and a carriage return.
- `>` – Waits for a `>`.
- `s\p9\c` – Sends an `s`, pauses, sends a `9` with no carriage return.

- `)-w\p\r\ds\p9\c-)` – Waits for a `)`. If `)` is not received, processes the string between the `-` characters as follows. Sends a `w`, pauses, sends a carriage return, delays, sends an `s`, pauses, sends a `9` without a carriage return, then waits for the `)`.
- `y\c` – Sends a `y` with no carriage return.
- `:` – Waits for a `:`.
- `\E\TP` – Enables echo checking. From this point on, whenever a character is transmitted, UUCP waits for the character to be received before proceeding, and then, sends the phone number. The `\T` means to take the phone number that is passed as an argument. The `\T` applies the `Dialcodes` translation and the modem function translation that is specified by field 2 of this entry. Then sends a `P` and a carriage return.
- `>` – Waits for a `>`.
- `9\c` – Sends a `9` without a newline.
- `OK` – Waits for the string `OK`.

UUCP Hardware Flow Control

You can also use the pseudo-send `STTY=value` string to set modem characteristics. For instance, `STTY=crtscts` enables outbound hardware flow control. `STTY=crtsxoff` enables inbound hardware flow control. `STTY=crtscts,crtsxoff` enables both outbound and inbound hardware flow control.

`STTY` accepts all the `stty` modes. See the `stty(1)` and `termio(7I)` man pages.

The following example would enable hardware flow control in a `Dialers` entry:

```
dsi =, -, "" \dA\pTE1V1X5Q0S2=255S12=255*E1*F3*M1*S1\r\c OK\r \EATDT\T\r\c
CONNECT\SEC STTY=crtscts
```

This pseudo-send string can also be used in entries in the `Systems` file.

UUCP Setting Parity

In some situations, you have to reset the parity because the system that you are calling checks port parity and drops the line if it is wrong. The expect-send couplet `P_ZERO` sets parity to zero:

```
foo =, -, "" P_ZERO "" \dA\pTE1V1X1Q0S2=255S12=255\r\c OK\r\EATDT\T\r\c CONNECT
```

In the same manner, `P_EVEN` sets parity to even (the default), `P_ODD` sets it to odd, and `P_ONE` sets it to one. This pseudo-send string can also be used in entries in the `Systems` file.

Other Basic UUCP Configuration Files

You can use files in this section in addition to the `Systems`, `Devices`, and `Dialers` file when doing basic UUCP configuration.

UUCP `/etc/uucp/Dialcodes` File

The `/etc/uucp/Dialcodes` file enables you to define dial-code abbreviations that can be used in the `Phone` field in the `/etc/uucp/Systems` file. You can use the `Dialcodes` files to provide additional information about a basic phone number that is used by several systems at the same site.

Each entry has the following format:

abbreviation dial-sequence

abbreviation represents the abbreviation that is used in the `Phone` field of the `Systems` file and *dial-sequence* represents the dial sequence that is passed to the dialer when that particular `Systems` file entry is accessed. The following table shows the correspondences between the two files.

TABLE 40-6 Correspondences Between `Dialcodes` and `Systems` Files

Field Names						
Dialcodes	<i>Abbreviation</i>	Dial-Sequence				
Systems	System-Name	Time	Type	Speed	<i>Phone</i>	Chat Script

The following table contains sample entries in a `Dialcodes` file.

TABLE 40-7 Entries in the `Dialcodes` File

Abbreviation	Dial-sequence
NY	1=212
jt	9+847

In the first row, `NY` is the abbreviation to appear in the `Phone` field of the `Systems` file. For example, the `Systems` file might have the following entry:

NY5551212

When `uucico` reads `NY` in the `Systems` file, it searches the `Dialcodes` file for `NY` and obtains the dialing sequence `1=212`. This is the dialing sequence needed for any phone call to New York City. This sequence includes the number `1`, an “equal sign” (=) meaning pause and wait for a secondary dial tone, and the area code `212`. `uucico` sends this information to the dialer, then returns to the `Systems` file for the remainder of the phone number, `5551212`.

The entry `jt 9=847-` would work with a `Phone` field such as `jt7867` in the `Systems` file. When `uucico` reads the entry that contains `jt7867` in the `Systems` file, it sends the sequence `9=847-7867` to the dialer, if the token in the dialer-token pair is `\T`.

UUCP /etc/uucp/Sysfiles File

The `/etc/uucp/Sysfiles` file lets you assign different files to be used by `uucp` and `cu` as `Systems`, `Devices`, and `Dialers` files. For more information on `cu`, see the `cu(1C)` man page. You can use `Sysfiles` for the following:

- Different `Systems` files so that requests for login services can be made to different addresses than `uucp` services.
- Different `Dialers` files so that you can assign different handshaking for `cu` and `uucp`.
- Multiple `Systems`, `Dialers`, and `Devices` files. The `Systems` file in particular can become large, making it more convenient to split it into several smaller files.

The format of the `Sysfiles` file is as follows:

```
service=w systems=x:x dialers=y:y devices=z:z
```

`w` represents `uucico`, `cu`, or both commands separated by a colon. `x` represents one or more files to be used as the `Systems` file, with each file name separated by a colon and read in the order that is presented. `y` represents one or more files to be used as the `Dialers` file. `z` is one or more files to be used as the `Devices` file.

Each file name is assumed to be relative to the `/etc/uucp` directory, unless a full path is given.

The following sample, `/etc/uucp/Sysfiles`, defines a local `Systems` file (`Local_Systems`) in addition to the standard `/etc/uucp/Systems` file:

```
service=uucico:cu systems=Systems :Local_Systems
```

When this entry is in `/etc/uucp/Sysfiles`, both `uucico` and `cu` first check in the standard `/etc/uucp/Systems`. If the system they are trying to call doesn't have an entry in that file, or if the entries in the file fail, then both commands check `/etc/uucp/Local_Systems`.

As specified in the previous entry, `cu` and `uucico` share the `Dialers` and `Devices` files.

When different Systems files are defined for uucico and cu services, your machine stores two different lists of Systems. You can print the uucico list by using the uuname command or the cu list by using the uuname -C command. The following is another example of the file, which shows that the alternate files are consulted first and the default files are consulted if necessary:

```
service=uucico systems=Systems.cico:Systems
dialers=Dialers.cico:Dialers \
devices=Devices.cico:Devices
service=cu systems=Systems.cu:Systems \
dialers=Dialers.cu:Dialers \
devices=Devices.cu:Devices
```

UUCP /etc/uucp/Sysname File

Every machine that uses UUCP must have an identifying name, often referred to as the *node name*. This is the name that appears in the remote machine's /etc/uucp/Systems file, along with the chat-script and other identifying information. Normally, UUCP uses the same node name as is returned by the uname -n command, which is also used by TCP/IP.

You can specify a UUCP node name independent of the TCP/IP host name by creating the /etc/uucp/Sysname file. The file has a one-line entry that contains the UUCP node name for your system.

UUCP /etc/uucp/Permissions File

The /etc/uucp/Permissions file specifies the permissions that remote computers have for login, file access, and command execution. Some options restrict the remote computer's ability to request files and its ability to receive files that are queued by the local machine. Another option is available that specifies the commands that a remote machine can execute on the local computer.

UUCP Structuring Entries

Each entry is a logical line, with physical lines terminated by a backslash (\) to indicate continuation. Entries are composed of options that are delimited by blank space. Each option is a name-value pair in the following format:

name=value

Values can be colon-separated lists. No blank space is allowed within an option assignment.

Comment lines begin with a pound sign (#), and they occupy the entire line up to a newline character. Blank lines are ignored (even within multiple-line entries).

The types of `Permissions` file entries are as follows:

- `LOGNAME` – Specifies the permissions that become effective when a remote computer logs in to (calls) your computer.

Note – When a remote machine calls you, its identity is questionable unless it has a unique login and verifiable password.

- `MACHINE` – Specifies permissions that become effective when your computer logs in to (calls) a remote computer.

`LOGNAME` entries contain a `LOGNAME` option and `MACHINE` entries contain a `MACHINE` option. One entry can contain both options.

UUCP Considerations

When using the `Permissions` file to restrict the level of access that is granted to remote computers, you should consider the following:

- All login IDs that are used by remote computers to log in for UUCP communications must appear in one and only one `LOGNAME` entry.
- Any site that is called with a name that does not appear in a `MACHINE` entry, has the following default permissions or restrictions:
 - Local send and receive requests are executed.
 - The remote computer can send files to your computer's `/var/spool/uucppublic` directory.
 - The commands that are sent by the remote computer for execution on your computer must be one of the default commands, usually `rmail`.

UUCP REQUEST Option

When a remote computer calls your computer and requests to receive a file, this request can be granted or denied. The `REQUEST` option specifies whether the remote computer can request to set up file transfers from your computer. The string `REQUEST=yes` specifies that the remote computer can request to transfer files from your computer. The string `REQUEST=no` specifies that the remote computer cannot request to receive files from your computer. `REQUEST=no`, the default value, is used if the `REQUEST` option is not specified. The `REQUEST` option can appear in either a `LOGNAME` entry (the remote computer calls you) or a `MACHINE` entry (you call remote computer).

UUCP SENDFILES Option

When a remote computer calls your computer and completes its work, it can attempt to retrieve work your computer has queued for it. The `SENDFILES` option specifies whether your computer can send the work that is queued for the remote computer.

The string `SENDFILES=yes` specifies that your computer can send the work that is queued for the remote computer if it is logged in as one of the names in the `LOGNAME` option. This string is *mandatory* if you have entered `Never` in the `Time` field of `/etc/uucp/Systems`. This designation sets up your local machine in passive mode, but it is not allowed to initiate a call to this particular remote computer. See “UUCP `/etc/uucp/Systems` File” on page 587 for more information.

The string `SENDFILES=call` specifies that files that are queued in your computer are sent only when your computer calls the remote computer. The `call` value is the default for the `SENDFILES` option. This option is only significant in `LOGNAME` entries because `MACHINE` entries apply when calls are sent to remote computers. If the option is used with a `MACHINE` entry, it is ignored.

UUCP MYNAME Option

This option enables you to designate a unique UUCP node name for your computer in addition to its TCP/IP host name, as returned by the `hostname` command. For instance, if you have unknowingly given your host the same name as that of some other system, you can set the `MYNAME` option of the `Permissions` file. Suppose that you want your organization to be known as `widget`. If all your modems are connected to a machine with the host name `gadget`, you can have an entry in `gadget's Permissions` file that reads as follows:

```
service=uucico systems=Systems.cico:Systems
  dialers=Dialers.cico:Dialers \
  devices=Devices.cico:Devices
service=cu systems=Systems.cu:Systems \
  dialers=Dialers.cu:Dialers \
  devices=Devices.cu:Devices
```

Now the system world can log in to the machine `gadget` as if it were logging in to `widget`. In order for machine world to know you also by the aliased name `widget` when you call it, you can have an entry that reads as follows:

```
MACHINE=world MYNAME=widget
```

You can also use the `MYNAME` option for testing purposes, as it allows your machine to call itself. However, because this option could be used to mask the real identity of a machine, you should use the `VALIDATE` option, as described in “UUCP `VALIDATE` Option” on page 610.

UUCP READ and WRITE Options

These options specify the various parts of the file system that `uucico` can read from or write to. You can designate `READ` and `WRITE` options with either `MACHINE` or `LOGNAME` entries.

The default for both the `READ` and `WRITE` options is the `uucppublic` directory, as shown in the following strings:

```
READ=/var/spool/uucppublic WRITE=/var/spool/uucppublic
```

The strings `READ=/` and `WRITE=/` specify permission to access any file that can be accessed by a local user with `Other` permissions.

The value of these entries is a colon-separated list of path names. The `READ` option is for requesting files, and the `WRITE` option is for depositing files. One of the values must be the prefix of any full path name of a file entering or exiting. To grant permission to deposit files in `/usr/news` as well as the public directory, use the following values with the `WRITE` option:

```
WRITE=/var/spool/uucppublic:/usr/news
```

If the `READ` and `WRITE` options are used, all path names must be specified because the path names are not added to the default list. For instance, if the `/usr/news` path name were the only path specified in a `WRITE` option, permission to deposit files in the public directory would be denied.

Be careful which directories you make accessible for reading and writing by remote systems. For example, the `/etc` directory contains many critical system files. Remote users should not have permission to deposit files in this directory.

UUCP NOREAD and NOWRITE Options

The `NOREAD` and `NOWRITE` options specify exceptions to the `READ` and `WRITE` options or defaults. The following entry permits reading any file except those files in the `/etc` directory (and its subdirectories—remember, these options are prefixes).

```
READ=/ NOREAD=/etc WRITE=/var/spool/uucppublic
```

This entry permits writing only to the default `/var/spool/uucppublic` directory. `NOWRITE` works in the same manner as the `NOREAD` option. You can use the `NOREAD` and `NOWRITE` options in both `LOGNAME` and `MACHINE` entries.

UUCP CALLBACK Option

You can use the `CALLBACK` option in `LOGNAME` entries to specify that no transaction occurs until the calling system is called back. The reasons to set up `CALLBACK` are as follows:

- For security purposes – If you call back a machine, you can be sure it is the right machine.
- For accounting purposes – If you are doing long data transmissions, you can choose the machine that is billed for the longer call.

The string `CALLBACK=yes` specifies that your computer must call back the remote computer before any file transfers can occur.

The default for the `CALLBACK` option is `CALLBACK=no`. If you set `CALLBACK` to `yes`, the permissions that affect the rest of the conversation must be specified in the `MACHINE` entry that corresponds to the caller. Do not specify these permissions in the `LOGNAME`, or in the `LOGNAME` entry that the remote machine might have set for your host.

Note – If two sites have the `CALLBACK` option set for each other, a conversation never is started.

UUCP COMMANDS Option



Caution – The `COMMANDS` option can compromise the security of your system. Use it with extreme care.

You can use the `COMMANDS` option in `MACHINE` entries to specify the commands that a remote computer can execute on your machine. The `uux` program generates remote execution requests and queues them to be transferred to the remote computer. Files and commands are sent to the target computer for remote execution. This is an exception to the rule that `MACHINE` entries apply only when your system calls out.

Note that `COMMANDS` is not used in a `LOGNAME` entry. `COMMANDS` in `MACHINE` entries defines command permissions, whether you call the remote system or it calls you.

The string `COMMANDS=rmail` specifies the default commands that a remote computer can execute on your computer. If a command string is used in a `MACHINE` entry, the default commands are overridden. For instance, the following entry overrides the `COMMAND` default so that the computers that are named `owl`, `raven`, `hawk`, and `dove` can now execute `rmail`, `rnews`, and `lp` on your computer.

```
MACHINE=owl:raven:hawk:dove COMMANDS=rmail:rnews:lp
```

In addition to the names as just specified, you can have full path names of commands. For example, the following entry specifies that command `rmail` uses the default search path.

```
COMMANDS=rmail:/usr/local/rnews:/usr/local/lp
```

The default search path for UUCP is `/bin` and `/usr/bin`. When the remote computer specifies `rnews` or `/usr/local/rnews` for the command to be executed, `/usr/local/rnews` is executed regardless of the default path. Likewise, `/usr/local/lp` is the `lp` command that is executed.

Including the `ALL` value in the list means that any command from the remote computers that are specified in the entry is executed. If you use this value, you give the remote computers full access to your machine.



Caution – This value allows far more access than normal users have. You should use this value only when both machines are at the same site, are closely connected, and the users are trusted.

Here is the string with the `ALL` value added:

```
COMMANDS=/usr/local/rnews:ALL:/usr/local/lp
```

This string illustrates two points:

- The `ALL` value can appear anywhere in the string.
- The path names that are specified for `rnews` and `lp` are used (instead of the default) if the requested command does not contain the full path names for `rnews` or `lp`.

You should use the `VALIDATE` option whenever you specify potentially dangerous commands, such as `cat` and `uucp` with the `COMMANDS` option. Any command that reads or writes files is potentially dangerous to local security when it is executed by the UUCP remote execution daemon (`uuxqt`).

UUCP VALIDATE Option

Use the `VALIDATE` option in conjunction with the `COMMANDS` option whenever you specify commands that are potentially dangerous to your machine's security. `VALIDATE` is merely an added level of security on top of the `COMMANDS` option, though it is a more secure way to open command access than `ALL`.

`VALIDATE` provides a certain degree of verification of the caller's identity by cross-checking the host name of a calling machine against the login name it uses. The following string ensures that if any machine other than `widget` or `gadget` tries to log in as `Uwidget`, the connection is refused.

```
LOGNAME=Uwidget VALIDATE=widget:gadget
```

The `VALIDATE` option requires privileged computers to have a unique login and password for UUCP transactions. An important aspect of this validation is that the login and password that are associated with this entry are protected. If an outsider obtains that information, that particular `VALIDATE` option can no longer be considered secure.

Carefully consider which remote computers you are granting privileged logins and passwords for UUCP transactions. Giving a remote computer a special login and password with file access and remote execution capability is like giving anyone on that computer a normal login and password on your computer. Therefore, if you cannot trust someone on the remote computer, do not provide that computer with a privileged login and password.

The following `LOGNAME` entry specifies that if one of the remote computers that claims to be `eagle`, `owl`, or `hawk` logs in on your computer, it must have used the login `uucpfriend`:

```
LOGNAME=uucpfriend VALIDATE=eagle:owl:hawk
```

If an outsider obtains the `uucpfriend` login and password, masquerading is easy.

But what does this entry have to do with the `COMMANDS` option, which appears only in `MACHINE` entries? This entry links the `MACHINE` entry (and `COMMANDS` option) with a `LOGNAME` entry that is associated with a privileged login. This link is needed because the execution daemon is not running while the remote computer is logged in. Actually, the link is an asynchronous process that does not know which computer sent the execution request. Therefore, the real question is, how does your computer know where the execution files came from?

Each remote computer has its own spool directory on your local machine. These spool directories have write permission that is given only to the UUCP programs. The execution files from the remote computer are put in its spool directory after being transferred to your computer. When the `uuxqt` daemon runs, it can use the spool directory name to find the `MACHINE` entry in the `Permissions` file and get the `COMMANDS` list. Or, if the computer name does not appear in the `Permissions` file, the default list is used.

This example shows the relationship between the `MACHINE` and `LOGNAME` entries:

```
MACHINE=eagle:owl:hawk REQUEST=yes \  
COMMANDS=rmail:/usr/local/rnews \  
READ=/ WRITE=/  
LOGNAME=uucpz VALIDATE=eagle:owl:hawk \  
REQUEST=yes SENDFILES=yes \  
READ=/ WRITE=/
```

The value in the `COMMANDS` option means that remote users can execute `rmail` and `/usr/local/rnews`.

In the first entry, you must assume that when you want to call one of the computers that is listed, you are really calling either eagle, owl, or hawk. Therefore, any files that are put into one of the eagle, owl, or hawk spool directories is put there by one of those computers. If a remote computer logs in and says that it is one of these three computers, its execution files are also put in the privileged spool directory. You therefore have to validate that the computer has the privileged login uucpz.

UUCP MACHINE Entry for OTHER

You might want to specify different option values for remote machines that are not mentioned in specific MACHINE entries. The need might arise when many computers are calling your host, and the command set changes from time to time. The name OTHER for the computer name is used for this entry as shown in this example:

```
MACHINE=OTHER \  
COMMANDS=rmail:rnews:/usr/local/Photo:/usr/local/xp
```

All other options available for the MACHINE entry can also be set for the computers that are not mentioned in other MACHINE entries.

Combining MACHINE and LOGNAME Entries for UUCP

You can combine MACHINE and LOGNAME entries into a single entry when the common options are the same. For example, the two sets of entries that follow share the same REQUEST, READ, and WRITE options:

```
MACHINE=eagle:owl:hawk REQUEST=yes \  
READ=/ WRITE=
```

and

```
LOGNAME=uupz REQUEST=yes SENDFILES=yes \  
READ=/ WRITE=
```

You can merge these entries, as shown:

```
MACHINE=eagle:owl:hawk REQUEST=yes \  
logname=uucpz SENDFILES=yes \  
READ=/ WRITE=
```

Combining MACHINE and LOGNAME entries makes the Permissions file more manageable and efficient.

UUCP Forwarding

When sending files through a series of machines, the intermediary machines must have the command `uucp` among their `COMMANDS` options. If you type the following command, the forwarding operation works only if machine `willow` permits machine `oak` to execute the `uucp` program.

```
% uucp sample.txt oak\!willow\!pine\!/usr/spool/uucppublic
```

The machine `oak` also must permit your machine to execute the `uucp` program. The machine `pine`, as the last machine designated, does not have to permit the `uucp` command because it is not doing any forwarding operations. Machines are not normally set up this way.

UUCP /etc/uucp/Poll File

The `/etc/uucp/Poll` file contains information for polling remote computers. Each entry in the `Poll` file contains the name of a remote computer to call, followed by a tab character or a space, and finally the hours the computer should be called. The format of entries in the `Poll` file are:

sys-name hour ...

For example, the entry

```
eagle 0 4 8 12 16 20
```

provides polling of computer `eagle` every four hours.

The `uudemon.poll` script processes the `Poll` file but does not actually perform the poll. The script merely sets up a polling work file (always named *C.file*) in the `spool` directory. The `uudemon.poll` script starts the scheduler, and the scheduler examines all work files in the `spool` directory.

UUCP /etc/uucp/Config File

The `/etc/uucp/Config` file enables you to override certain parameters manually. Each entry in the `Config` file has this format:

parameter=value

See the `Config` file that is provided with your system for a complete list of configurable parameter names.

The following `Config` entry sets the default protocol ordering to `Gge` and changes the `G` protocol defaults to 7 windows and 512-byte packets.

```
Protocol=G(7,512)ge
```

UUCP/etc/uucp/Grades File

The `/etc/uucp/Grades` file contains the definitions for the job grades that can be used to queue jobs to a remote computer. This file also contains the permissions for each job grade. Each entry in this file represents a definition of an administrator-defined job grade that lets users queue jobs.

Each entry in the `Grades` file has the following format:

User-job-grade System-job-grade Job-size Permit-type ID-list

Each entry contains fields that are separated by blank space. The last field in the entry is composed of subfields that are also separated by spaces. If an entry occupies more than one physical line, you can use a backslash to continue the entry onto the following line. Comment lines begin with a pound sign (#) and occupy the entire line. Blank lines are always ignored.

UUCP User-job-grade Field

This field contains an administrative-defined user job grade name of up to 64 characters.

UUCP System-job-grade Field

This field contains a single-character job grade to which *User-job-grade* is mapped. The valid list of characters is A-Z, a-z, with A having the highest priority and z the lowest.

Relationship Between User and System Job Grades

The user job grade can be bound to more than one system job grade. Note that the `Grades` file is searched sequentially for occurrences of a user job grade. Therefore, any multiple occurrences of a system job grade should be listed in compliance with the restriction on the maximum job size.

While no maximum number exists for the user job grades, the maximum number of system job grades that are allowed is 52. The reason is that more than one *User-job-grade* can be mapped to a *System-job-grade*, but each *User-job-grade* must be on a separate line in the file. Here is an example:

```
mail N Any User Any netnews N Any User Any
```

If this configuration is in a `Grades` file, these two *User-job-grade* fields share the same *System-job-grade*. Because the permissions for a *Job-grade* are associated with a *User-job-grade* and not a *System-job-grade*, two *User-job-grades* can share the same *System-job-grades* and have two different sets of permissions.

Default Grade

You can define the binding of a default *User-job-grade* to a system job grade. You must use the keyword `default` as user job grade in the *User-job-grade* field of the `Grades` file and the system job grade that it is bound to. The Restrictions and ID fields should be defined as `Any` so that any user and any size job can be queued to this grade. Here is an example:

```
default a Any User Any
```

If you do not define the default user job grade, the built-in default grade `Z` is used. Because the restriction field `default` is `Any`, multiple occurrences of the default grade are not checked.

UUCP Job-size Field

This field specifies the maximum job size that can be entered in the queue. *Job-size* is measured in bytes and can be a list of the options that are listed in the following table.

TABLE 40-8 Job-size Field

<i>nnnn</i>	Integer that specifies the maximum job size for this job grade
<i>nK</i>	Decimal number that represents the number of kilobytes (K is an abbreviation for kilobyte)
<i>nM</i>	Decimal number that represents the number of megabytes (M is an abbreviation for megabyte)
<i>Any</i>	Keyword that specifies that no maximum job size exists

Here are some examples:

- `5000` represents 5000 bytes
- `10K` represents 10 Kbytes
- `2M` represents 2 Mbytes

UUCP Permit-type Field

This field contains a keyword that denotes how to interpret the ID list. The following table lists the keywords and their meanings.

TABLE 40-9 Permit-type Field

Keyword	ID List Contents
User	Login names of users who are permitted to use this job grade
Non-user	Login names of users who are not permitted to use this job grade
Group	Group names whose members are permitted to use this group
Non-group	Group names whose members are not permitted to use this job grade

UUCP ID-list Field

This field contains a list of login names or group names that are to be permitted or denied queuing to this job grade. The list of names are separated by a blank space and terminated by a newline character. The keyword Any is used to denote that anyone is permitted to queue to this job grade.

Other UUCP Configuration Files

This section describes three less-frequently modified files that impact the use of UUCP facilities.

UUCP /etc/uucp/Devconfig File

The /etc/uucp/Devconfig file enables you to configure devices by service—uucp or cu. Devconfig entries define the STREAMS modules that are used for a particular device. They have the following format:

```
service=x device=y push=z[:z...]
```

x can be cu, uucico, or both services separated by a colon. *y* is the name of a network and must match an entry in the Devices file. *z* is replaced by the names of STREAMS modules in the order that they are to be pushed onto the Stream. Different modules and devices can be defined for cu and uucp services.

The following entries are for a STARLAN network and would most commonly be used in the file:

```
service=cu      device=STARLAN    push=ntty:tirdwr
service=uucico  device=STARLAN    push=ntty:tirdwr
```

This example pushes ntty, then tirdwr.

UUCP `/etc/uucp/Limits` File

The `/etc/uucp/Limits` file controls the maximum number of simultaneous `uucicos`, `uuxqts`, and `uuscheds` that are running in the uucp networking. In most situations, the default values are acceptable and no changes are needed. If you want to change them, however, use any text editor.

The format of the `Limits` file is as follows:

```
service=x max=y:
```

`x` can be `uucico`, `uuxqt` or `uusched`, and `y` is the limit that is permitted for that service. The fields can be in any order and in lowercase.

The following entries should most commonly be used in the `Limits` file:

```
service=uucico max=5
service=uuxqt max=5
service=uusched max=2
```

The example allows five `uucicos`, five `uuxqts`, and two `uuscheds` running on your machine.

UUCP `remote.unknown` File

The other file that affects the use of communication facilities is the `remote.unknown` file. This file is a binary program that executes when a machine that is not found in any of the `Systems` files starts a conversation. This program logs the conversation attempt and drops the connection.



Caution – If you change the permissions of the `remote.unknown` file so that it cannot execute, your system accepts connections from any system.

This program executes when a machine that is not in any of the `Systems` starts a conversation. The program logs the conversation attempt but fails to make a connection. If you change the permissions of this file so that it cannot execute (`chmod 000 remote.unknown`), your system accepts any conversation requests. This is not a trivial change, and you should have good reasons for making it.

UUCP Administrative Files

The UUCP administrative files are described next. These files are created in spool directories to lock devices, hold temporary data, or keep information about remote transfers or executions.

- *Temporary data files (TM)* – These data files are created by UUCP processes under the spool directory `/var/spool/uucp/x` when a file is received from another computer. The directory `x` has the same name as the remote computer that is sending the file. The names of the temporary data files have the following format:

`TM.pid.ddd`

`pid` is a process ID and `ddd` is a sequential three-digit number that starts at 0.

When the entire file is received, the `TM.pid.ddd` file is moved to the path name that is specified in the `C.sysnxxxx` file (discussed subsequently) that caused the transmission. If processing is abnormally terminated, the `TM.pid.ddd` file can remain in the `x` directory. These files should be automatically removed by `uucleanup`.

- *Lock files (LCK)* – Lock files are created in the `/var/spool/locks` directory for each device in use. Lock files prevent duplicate conversations and multiple attempts to use the same calling device. The following table shows the different types of UUCP lock files.

TABLE 40-10 UUCP Lock Files

File Name	Description
<code>LCK.sys</code>	<code>sys</code> represents the name of the computer that is using the file
<code>LCK.dev</code>	<code>dev</code> represents the name of a device that is using the file
<code>LCK.LOG</code>	<code>LOG</code> represents a locked UUCP log file

These files can remain in the spool directory if the communications link is unexpectedly dropped (usually on computer crashes). The lock file is ignored (removed) after the parent process is no longer active. The lock file contains the process ID of the process that created the lock.

- *Work file (C.)* – Work files are created in a spool directory when work (file transfers or remote command executions) has been queued for a remote computer. The names of work files have the following format:

`C.sysnxxxx`

`sys` is the name of the remote computer, `n` is the ASCII character that represents the grade (priority) of the work, and `xxxx` is the four-digit job sequence number that is assigned by UUCP. Work files contain the following information:

- Full path name of the file to be sent or requested.
- Full path name of the destination or user or file name.
- User login name.
- List of options.
- Name of associated data files in the spool directory. If the `uucp -C` or `uuto -p` option was specified, a dummy name (`D. 0`) is used.
- Mode bits of the source file.
- Remote user's login name to be notified on completion of the transfer.
- *Data file(D.)* – Data files are created when you specify on the command line to copy the source file to the spool directory. The names of data files have the following format:

`D.systmxxxxyyy` – `systm` is the first five characters in the name of the remote computer, `xxxx` is a four-digit job sequence number assigned by `uucp`. The four-digit job sequence number can be followed by a subsequence number. `yyy` is used when several `D.` files are created for a work (`C.`) file.
- *X. (execute file)* – Execute files are created in the spool directory prior to remote command executions. The names of execute files have the following format:

`X.sysnxxxx`

`sys` is the name of the remote computer. `n` is the character that represents the grade (priority) of the work. `xxxx` is a four-digit sequence number that is assigned by UUCP. Execute files contain the following information:

 - Requester's login and computer name
 - Names of files that are required for execution
 - Input to be used as the standard input to the command string
 - Computer and file name to receive standard output from the command execution
 - Command string
 - Option lines for return status requests

UUCP Error Messages

This section lists the error messages that are associated with UUCP.

UUCP ASSERT Error Messages

The following table lists ASSERT error messages.

TABLE 40–11 ASSERT Error Messages

Error Message	Description or Action
CAN'T OPEN	An <code>open()</code> or <code>fopen()</code> failed.
CAN'T WRITE	A <code>write()</code> , <code>fwrite()</code> , <code>fprint()</code> , or similar command, failed.
CAN'T READ	A <code>read()</code> , <code>fgets()</code> , or similar command failed.
CAN'T CREATE	A <code>creat()</code> call failed.
CAN'T ALLOCATE	A dynamic allocation failed.
CAN'T LOCK	An attempt to make a LCK (lock) file failed. In some situations, this is a fatal error.
CAN'T STAT	A <code>stat()</code> call failed.
CAN'T CHMOD	A <code>chmod()</code> call failed.
CAN'T LINK	A <code>link()</code> call failed.
CAN'T CHDIR	A <code>chdir()</code> call failed.
CAN'T UNLINK	An <code>unlink()</code> call failed.
WRONG ROLE	This is an internal logic problem.
CAN'T MOVE TO CORRUPTDIR	An attempt to move some bad C. or X. files to the <code>/var/spool/uucp/</code> . Corrupt directory failed. The directory is probably missing or has wrong modes or owner.
CAN'T CLOSE	A <code>close()</code> or <code>fclose()</code> call failed.
FILE EXISTS	The creation of a C. or D. file is attempted, but the file exists. This error occurs when a problem arises with the sequence file access. Usually indicates a software error.
NO uucp SERVICE NUMBER	A TCP/IP call is attempted, but no entry is in <code>/etc/services</code> for UUCP.
BAD UID	The user ID is not in the password database. Check name service configuration.
BAD LOGIN_UID	Same as previous description.
BAD LINE	A bad line is in the <code>Devices</code> file. Not enough arguments on one or more lines.
SYSLST OVERFLOW	An internal table in <code>gename.c</code> overflowed. A single job attempted to talk to more than 30 systems.
TOO MANY SAVED C FILES	Same as previous description.
RETURN FROM fixline ioctl	An <code>ioctl(2)</code> , which should never fail, failed. A system driver problem has occurred.
BAD SPEED	A bad line speed appears in the <code>Devices</code> or <code>Systems</code> file (Class or Speed field).
BAD OPTION	A bad line or option is in the <code>Permissions</code> file. It must be fixed immediately.
PKCGET READ	The remote machine probably hung up. No action need be taken.

TABLE 40–11 ASSERT Error Messages (Continued)

Error Message	Description or Action
PKXSTART	The remote machine aborted in a nonrecoverable way. This error can usually be ignored.
TOO MANY LOCKS	An internal problem has occurred. Contact your system vendor.
XMV ERROR	A problem with some file or directory has occurred. The spool directory is the probable cause, as the modes of the destinations were supposed to be checked before this process was attempted.
CAN'T FORK	An attempt to make a <code>fork</code> and <code>exec</code> failed. The current job should not be lost but will be attempted later (<code>uuxqt</code>). No action is needed.

UUCP STATUS Error Messages

The following table is a list of the most common STATUS error messages.

TABLE 40–12 UUCP STATUS Messages

Error Message	Description/Action
OK	Status is okay.
NO DEVICES AVAILABLE	Currently no device is available for the call. Check whether a valid device is in the <code>Devices</code> file for the particular system. Check the <code>Systems</code> file for the device to be used to call the system.
WRONG TIME TO CALL	A call was placed to the system at a time other than what is specified in the <code>Systems</code> file.
TALKING	Self-explanatory.
LOGIN FAILED	The login for the particular machine failed. The cause could be a wrong login or password, wrong number, a slow machine, or failure in executing the <code>Dialer-Token-Pairs</code> script.
CONVERSATION FAILED	The conversation failed after successful startup. This error usually means that one side went down, the program aborted, or the line (link) was dropped.
DIAL FAILED	The remote machine never answered. The cause could be a bad dialer or the wrong phone number.
BAD LOGIN/MACHINE COMBINATION	The machine called with a login/machine name that does not agree with the <code>Permissions</code> file. This error could be an attempt to masquerade.
DEVICE LOCKED	The calling device to be used is currently locked and in use by another process.

TABLE 40–12 UUCP STATUS Messages *(Continued)*

Error Message	Description/Action
ASSERT ERROR	An ASSERT error occurred. Check the <code>/var/uucp/.Admin/errors</code> file for the error message and refer to the section “UUCP ASSERT Error Messages” on page 619.
SYSTEM NOT IN Systems FILE	The system is not in the <code>Systems</code> file.
CAN'T ACCESS DEVICE	The device tried does not exist or the modes are wrong. Check the appropriate entries in the <code>Systems</code> and <code>Devices</code> files.
DEVICE FAILED	The device could not be opened.
WRONG MACHINE NAME	The called machine is reporting a different name than expected.
CALLBACK REQUIRED	The called machine requires that it call your machine.
REMOTE HAS A LCK FILE FOR ME	The remote machine has a LCK file for your machine. The remote machine could be trying to call your machine. If it has an older version of UUCP, the process that was talking to your machine might have failed, leaving the LCK file. If the remote machine has the new version of UUCP and is not communicating with your machine, the process that has a LCK file is hung.
REMOTE DOES NOT KNOW ME	The remote machine does not have the node name of your machine in its <code>Systems</code> file.
REMOTE REJECT AFTER LOGIN	The login that was used by your machine to log in does not agree with what the remote machine was expecting.
REMOTE REJECT, UNKNOWN MESSAGE	The remote machine rejected the communication with your machine for an unknown reason. The remote machine might not be running a standard version of UUCP.
STARTUP FAILED	Login succeeded, but initial handshake failed.
CALLER SCRIPT FAILED	This error is usually the same as DIAL FAILED. However, if this error occurs often, suspect the caller script in the <code>Dialers</code> file. Use <code>Uutry</code> to check.

UUCP Numerical Error Messages

The following table lists the exit code numbers of error status messages that are produced by the `/usr/include/sysexits.h` file. Not all are currently used by uucp.

TABLE 40–13 UUCP Error Messages by Number

Message Number	Description	Meaning
64	Base Value for Error Messages	Error messages begin at this value.

TABLE 40–13 UUCP Error Messages by Number *(Continued)*

Message Number	Description	Meaning
64	Command-Line Usage Error	The command was used incorrectly, for example, with the wrong number of arguments, a bad flag, or a bad syntax.
65	Data Format Error	The input data was incorrect in some way. This data format should only be used for user's data and not system files.
66	Cannot Open Input	An input file (not a system file) did not exist, or was not readable. This problem could also include errors like "No message" to a mailer.
67	Address Unknown	The user that was specified did not exist. This error might be used for mail addresses or remote logins.
68	Host Name Unknown	The host did not exist. This error is used in mail addresses or network requests.
69	Service Unavailable	A service is unavailable. This error can occur if a support program or file does not exist. This message also can be a catchall message when something doesn't work and you don't know why.
70	Internal Software Error	An internal software error has been detected. This error should be limited to non-operating system-related errors, if possible.
71	System Error	An operating system error has been detected. This error is intended to be used for conditions like "cannot fork", "cannot create pipe." For instance, this error includes a <code>getuid</code> return of a user who does not exist in the <code>passwd</code> file.
72	Critical OS File Missing	Some system file such as <code>/etc/passwd</code> or <code>/var/admin/utmpx</code> does not exist, cannot be opened, or has some error such as syntax error.
73	Can't Create Output File	A user-specified output file cannot be created.
74	Input/Output Error	An error occurred while doing I/O on some file.
75	Temporary Failure. User is invited to retry	Temporary failure, indicating something that is not really an error. In <code>sendmail</code> , this means that a mailer, for example, could not create a connection, and the request should be reattempted later.
76	Remote Error in Protocol	The remote system returned something that was "not possible" during a protocol exchange.
77	Permission Denied	You do not have sufficient permission to perform the operation. This message is not intended for file system problems, which should use <code>NOINPUT</code> or <code>CANTCREAT</code> , but rather for higher-level permissions. For example, <code>kre</code> uses this message to restrict students who can send mail to.
78	Configuration Error	The system detected an error in the configuration.
79	Entry Not Found	Entry not found.
79	Maximum Listed Value	Highest value for error messages.

Composed September 5, 2002

CHAPTER **41**

Working With Remote Systems Topics

This section provides instructions for administering an FTP Server and for accessing remote systems in the Solaris environment. The section contains these chapters.

Chapter 43	Step-by-step instructions for administering the FTP Server.
Chapter 44	Step-by-step instructions for accessing remote files.

Composed September 5, 2002

Working With Remote Systems (Overview)

This section includes information on working with remote files.

What Is the FTP Server?

The FTP Server is based on wu-ftp. Originally developed by Washington University in Saint Louis, wu-ftp is widely used for distribution of bulk data over the Internet and is the preferred standard for large FTP sites. For information on the licensing terms, refer to the materials that are incorporated at `/var/sadm/pkg/SUNWftpu/install/copyright`.

What Is a Remote System?

For the purpose of this chapter, a remote system is a workstation or server that is connected to the local system with any type of physical network and configured for TCP/IP communication.

On systems running the Solaris 9 release, TCP/IP configuration is established automatically during startup. For more information, see *System Administration Guide: IP Services*.

What's New for the Solaris 9 FTP Server?

The FTP Server is compatible with Solaris 8 FTP software, yet offers new capability with improved performance for Solaris 9 users.

TABLE 42-1 What's New for the Solaris 9 FTP Server

Feature	Description	For Information
User classification by type and location	Permits you to define a class of users, based on type and address	"How to Define FTP Server Classes" on page 633
Limits per class	Controls the number of users from a certain class who are allowed simultaneous login, based on limits that are set in the <code>ftppaccess</code> file	"How to Set User Login Limits" on page 634
System-wide and directory-related messages	Displays the messages that you specify for particular events	"How to Create Messages to Be Sent to Users" on page 642
Upload permissions per directory	Allows you to control uploads to the FTP Server, including file and directory creation and permissions	"How to Control Uploads to the FTP Server" on page 646
File name filter	Enables you to specify which characters, in what sequence, are acceptable in the name of an uploaded file	"How to Control Uploads to the FTP Server" on page 646
Virtual host support	Permits you to configure the FTP server to support multiple domains on a single machine	"How to Enable Complete Virtual Hosting" on page 650
Command logging	Allows logging of commands that are executed by real, guest, and anonymous FTP users	"How to Check the Commands Executed by FTP Users" on page 655
Transfer logging	Allows logging of transfers for real, guest, and anonymous FTP users	<code>ftppaccess(4)</code> , <code>xferlog(4)</code> , in <code>ftpd(1M)</code>
As-needed compression and archiving	Allows as-needed compression and archiving by using conversions that are specified in the <code>ftpconversions</code> file	<code>ftpconversions(4)</code> , <code>ftppaccess(4)</code>

Note – The Solaris 8 `/etc/default/ftpd` is not supported in the Solaris 9 release. During upgrade, `BANNER` and `UMASK` entries are converted to their `wu-ftpd` equivalents. However, the system administrator might need to manually convert some `BANNER` lines for the equivalent `ftppass` greeting capability. For further information, see `ftppass(4)`.

Note – The sublogin feature that is provided by the Solaris 8 FTP Server is not supported by the Solaris 9 FTP Server.

Composed September 5, 2002

Administering the FTP Server (Tasks)

This chapter includes tasks that are described in the following table to set up and administer an FTP server.

TABLE 43-1 Task Map: Administering the FTP Server

Task	Description	For Instructions
Configure access to the FTP server	Use the <code>ftppass</code> , <code>ftpusers</code> , and the <code>ftphosts</code> files in the <code>/etc/ftpd</code> directory to establish or restrict access to the FTP server.	"How to Set User Login Limits" on page 634 "How to Control the Number of Invalid Login Attempts" on page 635 "How to Disallow FTP Server Access to Particular Users" on page 636 "How to Restrict Access to the Default FTP Server" on page 637 "How to Define FTP Server Classes" on page 633
Set up FTP server logins	Establish login accounts for real, guest and anonymous users.	"How to Set Up Real FTP Users" on page 638 "How to Set Up Guest FTP Users" on page 639 "How to Set Up Anonymous FTP Users" on page 640 "How to Create the <code>/etc/shells</code> file" on page 640
Customize message files	Edit the <code>/etc/ftpd/ftppass</code> file to configure the FTP server to return messages to the FTP client related to specific events.	"How to Customize Message Files" on page 641 "How to Create Messages to Be Sent to Users" on page 642 "How to Configure the README Option" on page 643

TABLE 43-1 Task Map: Administering the FTP Server *(Continued)*

Task	Description	For Instructions
Configure access to files on the FTP server	Use the <code>/etc/ftpd/ftpaccess</code> file to specify classes of users who are allowed to execute certain commands or to download and upload files to the FTP server.	“How to Configure DA Discovery for Dial-up Networks” on page 270 “Controlling Uploads and Downloads on the FTP Server” on page 645
Enable limited or complete virtual hosting	Use the <code>/etc/ftpd/ftpaccess</code> file to configure the FTP server to support multiple domains on the same machine.	“How to Enable Limited Virtual Hosting” on page 649 “How to Enable Complete Virtual Hosting” on page 650
Start the FTP server	Edit the <code>/etc/inet/inetd.conf</code> file to start the FTP server in <code>nowait</code> or <code>standalone</code> mode.	“How to Start an FTP Server From <code>inetd.conf</code> ” on page 652 “How to Start a Standalone FTP Server” on page 653
Shut down the FTP server	Use the <code>/etc/ftpd/ftpaccess</code> file and run the <code>ftpsht</code> to shut down the FTP server.	“Shutting Down the FTP Server” on page 653
Troubleshoot some common FTP server problems	Check <code>syslogd</code> and use greeting text and log commands to debug problems on the FTP server.	“How to Check <code>syslogd</code> for FTP Server Messages” on page 655 “How to Use greeting text to Verify <code>ftpaccess</code> ” on page 655 “How to Check the Commands Executed by FTP Users” on page 655

Controlling FTP Server Access

You can use the following configuration files in the `/etc/ftpd` directory to control access to the FTP server.

- `ftpusers` is used to list users who are denied access to the FTP server.
- `ftphosts` is used to allow or deny login from various hosts to various accounts on the FTP server.
- `ftpaccess` is the main FTP configuration file. The FTP server only reads the `/etc/ftpd/ftpaccess` file if called with the `-a` option. When the `ftpaccess` file is used, all users must be members of a class to be allowed access to the FTP server. You can specify many `ftpaccess` directives that apply only to a particular class.

For further information, see `ftpusers(4)`, `ftphosts(4)`, and `ftppass(4)`

Note – In all FTP server configuration files, lines beginning with # signs are treated as comments.

▼ How to Define FTP Server Classes

To log in to the FTP server, users must be members of a class when the `ftppass` file is used. To add the `class` directive to the `ftppass` file, you specify the *class* name, *typelist* of users who are permitted access from a particular host.

1. **Become superuser.**
2. **Add entries for anonymous, guest, and real users in the `ftppass` file.**

```
class class typelist addrglob[addrglob...]
```

<code>class</code>	Keyword that is used to define FTP users.
<code>class</code>	A name that is defined by the <code>class</code> keyword. Each login is compared against a list of defined classes. The logged in user is considered a member of the first class matched.
<code>typelist</code>	A comma-separated list of the keywords that match the three types of users: <code>anonymous</code> , <code>guest</code> , and <code>real</code> .
<code>addrglob</code>	A globbed domain name or a globbed numeric address. The <i>addrglob</i> can also be the name of a file, starting with a slash ('/'), which contains additional address globs: <code>address:netmask</code> or <code>address/cidr</code>. Here are some examples of globbed addresses: ■ Numeric IPv4 address: <code>10.1.2.3</code> ■ Globbed domain name <code>*.provider.com</code> ■ Globbed numeric IPv4 address <code>10.1.2.*</code> ■ Numeric IPv4 address:netmask <code>10.1.2.0:255.255.255.0</code> ■ Numeric IPv4 address/CIDR <code>10.1.2.0/24</code> ■ Numeric IPv6 address: <code>2000::56:789:21ff:fe8f:ba98</code> ■ Numeric IPv6 address/CIDR: <code>2000::56:789:21ff:fe8f:ba98/120</code>

Example—Defining FTP Server Classes

```
class local real,guest,anonymous *.provider.com
class remote real,guest,anonymous *
```

The previous example defines the `local` class as any user of the type `real`, `guest`, or `anonymous` who logs in from `*.provider.com`. The last line defines `remote` as any user who logs in from anywhere other than `*.provider.com`.

▼ How to Set User Login Limits

You can limit the number of simultaneous logins by users of a certain class with directives that are set in the `ftppaccess` file. Each login limit contains the name of a class, a UUCP-style days-of-week list, and a message file to display if the limit is exceeded.

To set user login limits, follow the steps in the next procedure.

1. **Become superuser.**
2. **Add the following entries to the `ftppaccess` file:**

```
limit class n times [message_file]
```

`limit`

Keyword that is used to restrict simultaneous logins by the specified number of users of a defined class at certain connection times.

`class`

A name that is defined by the `class` keyword. Each login is compared against a list of defined classes. The logged-in user is considered a member of the first class matched.

`n`

Number of users.

`times`

Day-of-week and time-of-day when the class can connect. Use `Any` for any day.

`message_file`

Message file that is displayed if a user is denied access.

Example—Setting User Login Limits

```
limit anon 50 Wk0800-1800 /etc/ftpd/ftpmsg.deny
limit anon 100 Any /etc/ftpd/ftpmsg.deny
limit guest 100 Any /etc/ftpd/ftpmsg.deny
```

The first line of the preceding example shows a limit of 50 simultaneous logins that are allowed to users of class `anon` during weekly work hours. The second line limits `anon` users to 100 simultaneous logins outside of working hours. The last line shows a limit of 100 `guest` logins that are allowed at any time. For information on how to specify day and time parameters, see `ftppaccess(4)`.

The example further indicates that the content of the file `/etc/ftpd/ftpmsg.deny` is returned when a specified login limit is reached, assuming `ftpmsg.deny` exists. For information on using the `/usr/sbin/ftpcount` command to view the number and login limit for each class of user who is logged in at a particular time, see `ftpcount(1)`.

Users are allowed login to the FTP server unless a specified limit is reached. Anonymous users are logged in as the user `ftp`. Real users are logged in as themselves, and guests are logged in as real users with a `chroot` environment to limit access privileges.

For information on using the `/usr/sbin/ftpwho` command to check the identities of the users logged into the FTP server, see `ftpwho(1)`.

▼ How to Control the Number of Invalid Login Attempts

If a login to the FTP server fails because of a problem such as misspelling required information, login is usually repeated. The user is allowed a specific number of consecutive login attempts before a message is logged to the `syslog` file. At that point, the user is disconnected. You can set a failure limit on the number of login attempts by following steps in the next procedure.

1. **Become superuser.**
2. **Add the following entries to the `ftppass` file.**

```
loginfails n
```

loginfails

Keyword that is used to assign the number of login failures that are permitted before the FTP connection is terminated

n

Number of times a login can fail

Example—Controlling the Number of Invalid Login Attempts

```
loginfails 10
```

The preceding example states that the user is disconnected from the FTP server after 10 failed login attempts.

▼ How to Disallow FTP Server Access to Particular Users

The `/etc/ftpd/ftpusers` file lists names of users who are not allowed to log in to the FTP server. When login is attempted, the FTP server checks the `/etc/ftpd/ftpusers` file to determine whether the user should be denied access. If the user's name is not found in that file, the server then searches the `/etc/ftpusers` file.

If the user's name is matched in `/etc/ftpusers`, a `syslogd` message is written with a statement that the match was found in a deprecated file. The message also recommends the use of `/etc/ftpd/ftpusers` instead of `/etc/ftpusers`.

Note – Support for the `/etc/ftpusers` file has been deprecated in this release. If the `/etc/ftpusers` file exists when the FTP server is installed, the file is moved to `/etc/ftpd/ftpusers`.

For additional information, see `syslogd(1M)`, `in.ftpd(1M)`, and `ftpusers(4)`

1. **Become superuser.**
2. **Add entries to the `/etc/ftpd/ftpusers` file for users who are not allowed to log in to the FTP server.**

Example—How to Disallow FTP Server Access

```
root
daemon
bin
sys
adm
lp
uccp
nuucp
listen
nobody
noaccess
nobody4
```

The previous example lists the typical entries in the `ftpusers` file. User names match entries in the `/etc/passwd`. The list generally includes the superuser `root` and other administrative and system application identities.

The `root` entry is included in the `ftpusers` file as a security measure. The default security policy is to disallow remote logins for `root`. The policy is also followed for the default value that is set as the `CONSOLE` entry in the `/etc/default/loginfile`. See `login(1)`.

▼ How to Restrict Access to the Default FTP Server

In addition to the controls mentioned previously, you can add explicit statements to the `ftppaccess` file to restrict access to the FTP server.

1. **Become superuser.**
2. **Add the following entries to the `ftppaccess` file.**
 - a. **By default, all users are allowed access to the default (non-virtual) FTP server. To deny access for specific users (other than `anonymous`), add the following entry:**

```
defaultserver deny username [username...]
```

`defaultserver`

Keyword that is used to identify the non-virtual server to which access can be denied or allowed

`username`

Login name of a user with restricted access to the `defaultserver`

- b. **To allow access for users who are not listed on the `deny` line, add the following line:**

```
defaultserver allow username [username...]
```

- c. **To prevent access by anonymous users, add the entry:**

```
defaultserver private
```

Example—Restricting Access to the Default FTP Server

```
defaultserver deny *  
defaultserver allow username
```

The previous example states that the FTP server denies access to all users except `anon` users and those users who are listed on the `allow` line.

You can also use the `ftphosts` file to deny access to particular login accounts from various hosts. See `ftphosts(4)` for additional information.

Setting Up FTP Server Logins

To access an FTP server, you must first log in. The FTP server supports three types of user login accounts for *real*, *guests*, and *anonymous* users.

- *Real* users have accounts that allow them to establish terminal sessions on systems that run the FTP server. Subject to directory and file access permissions, the entire disk structure is visible to real users.
- *Guest* users also need accounts to log in to the FTP server. Each guest account is set up with a user name and password. Functioning login shells are not assigned to guests to prevent users from establishing terminal sessions. At login, the FTP server performs a `chroot(2)` operation to restrict a guest's view of the server's disk structure.

Note – Login shells for real and guest users must be listed in the `/etc/shells` file to allow access to the FTP server.

- *Anonymous* users log in to the FTP server by using the either `ftp` or `anonymous` as a user name. By convention, anonymous users supply an email address when prompted for a password.

At login, the FTP server performs a `chroot(2)` operation that restricts the anonymous user's view of the server's disk structure. A single file area is shared by all anonymous users, unlike the separate areas that can be created for each guest user.

Real and guest users log in by using individual accounts with passwords that are known only to one person. Anonymous users log in to a well-known account that is potentially available to anyone. Most large-scale file distribution is created by using the anonymous account.

▼ How to Set Up Real FTP Users

To enable access for real users to the FTP server, follow these instructions:

1. **Verify that the user has an account that is set up with a user name and password that can be used to establish a terminal session.**

For more information, see “Managing User Accounts and Groups (Overview)” in the *System Administration Guide, Volume 1*.

2. **Confirm that the real user is a member of a class in the `ftppaccess` file.**

For information on the user classes that are defined in the `ftppaccess` file, see “How to Define FTP Server Classes” on page 633.

3. Verify that the user's login shell is listed in the `/etc/shells` file.

▼ How to Set Up Guest FTP Users

The `ftpconfig` script is used to copy all necessary system files to the home directory. When the guest user and the guest's home directory already exist, the `ftpconfig` script updates the area with the current system files.

For more information, see `ftpconfig(1M)`

Note – Unlike the user name (anonymous or `ftp`) that is set for anonymous users, user names for FTP guests are not fixed. Any name that would work as a real user name can be selected.

To enable access by a guest user to the FTP server, do the following:

1. Use the `useradd` script to create a guest user account with a login shell of `/bin/true` and a home directory of `/root_dir/.home_dir`.

For more information, see `useradd(1M)` and “Managing Use Accounts and Groups (Overview)” in the *System Administration Guide, Volume 1*.

Note – In this procedure, `/home/guests/.guest1` is used as the home directory name for a user who is called `guest1`.

```
# /usr/sbin/useradd -m -c "Guest FTP" -d \  
/home/guests/.guest1 -s /bin/true guest1
```

2. Assign a password to the guest account.
3. Add a `guestuser` entry to the `ftpaccess` file.

```
guestuser guest1
```

Note – You can also use the `guestgroup` capability in the `ftpaccess` file to specify guest users. The `guest-root` capability in `ftpaccess` eliminates the need for the `./` in the guest user's home directory path.

4. Confirm that the guest user is a member of a class in the `ftpaccess` file. See “How to Define FTP Server Classes” on page 633 for further information.
5. Use the `ftpconfig` script to create the required files in the `chroot` area.

```
/usr/sbin/ftpconfig -d /home/guests
```

6. **Confirm that `/bin/true` is listed in the `/etc/shells` file. See “How to Create the `/etc/shells` file” on page 640.**

Example—Setting Up a Guest FTP Server

In this example, the FTP area is set up in the `/home/guests` directory.

```
# /usr/sbin/ftpconfig -d /home/guests
Updating directory /home/guests
```

▼ How to Set Up Anonymous FTP Users

The `ftpconfig` script creates the anonymous user account and populates the home directory with the required files.

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

To enable access by an anonymous user to the FTP server, follow these instructions:

1. **Use the `ftpconfig` script to create the anonymous user account.**

```
# /usr/sbin/ftpconfig anonymous-ftp-directory
```
2. **Confirm that the anonymous user is assigned to a class in the `ftppass` file.**
See “How to Define FTP Server Classes” on page 633 for further information.

Example—Setting Up Anonymous FTP Users

In this example, the FTP area is set up in the `/home/ftp` directory.

```
# /usr/sbin/ftpconfig /home/ftp
Creating user ftp
Updating directory /home/ftp
```

▼ How to Create the `/etc/shells` file

1. **Become superuser.**
2. **Create the `/etc/shells` file.**
3. **Edit `/etc/shells`. Add the full path to each shell on a single line.**

Example—Creating the `/etc/shells` file

The following is an example of an `/etc/shells` file with a `/bin/true` listed for FTP guest users:


```
/sbin/sh
/bin/csh
/bin/jsh
/bin/ksh
/bin/remsh
/bin/rksh
/bin/rsh
/bin/sh
/usr/bin/csh
/usr/bin/ksh
/usr/bin/bash
/usr/bin/tcsh
/usr/bin/zsh
/bin/true
```

Customizing Message Files

You can configure the FTP server to return messages that are related to specific events to the FTP client. A welcome message might be set to display when a user logs in to the FTP server. Another message could appear when the user makes a directory change.

In addition to plain text, message files can contain one or more *magic cookies*. A magic cookie is composed of a % (percent sign), followed by a single character. When you embed a cookie in message text, information that is associated with the cookie appears on screen at the point the message file is called.

For example, message text might contain the cookie %L:

```
Welcome to %L!
```

When the message is displayed, the magic cookie %L is replaced with the name of the server as defined by the `hostname` statement in the `ftppaccess` file. For a complete list of supported message cookies, see `ftppaccess(4)`.

Note – If the host name is not defined in the `ftppaccess` file, the default host name for the local machine is used.

▼ How to Customize Message Files

1. **Become superuser.**
2. **Edit your message file to include magic cookies as appropriate.**
See `ftppaccess(4)` for a list of cookies you can use.

Example—Customizing Message Files

The following is an example of a message file that includes magic cookies:

```
Welcome to %L -- local time is %T.

You are number %N out of a maximum of %M.
All transfers are logged.

If your FTP client crashes or hangs shortly after login
please try
using a dash (-) as the first character of your password.
This will
turn off the informational messages that may be confusing
your FTP
client.

Please send any comments to %E.
```

▼ How to Create Messages to Be Sent to Users

After the user is logged in, system-related or application-related messages are displayed on screen. The `ftppaccess` file lists the events that trigger associated message statements.

1. **Become superuser.**
2. **Add the following entries to the `ftppaccess` file:**

```
message message_file [when [class ...]]
```

message

Keyword that is used to specify the message file to be displayed when a user logs in or executes the command to change the working directory.

message_file

Name of the message file to be displayed.

when

Parameter that is set as `login` or `cwd=dir`. See the following example.

class

The `class` specification allows the message to be displayed only to members of a particular class.

Example: How to Create Messages to Be Sent to Users

```
message /etc/ftpd/Welcome login anon guest
message .message cwd=*
```

The preceding example states that the file `/etc/ftpd/Welcome` is displayed at login for users of the class `anon` or `guest`. The second line states that the `.message` file in the current working directory is displayed for all users.

Message files are created relative to the `chroot` directory for guest and anonymous users.

▼ How to Configure the README Option

The first time a directory is visited, README files can be listed. To configure the README option, add the following entries to the `ftppaccess` file.

1. **Become superuser.**
2. **Add the following entries to the `ftppaccess` file.**

```
readme message_file [when [class...]]
```

readme

Keyword that is used to specify a message file to be checked when a user logs in or changes the working directory. If the message file exists, the user is notified and is given the date the file was modified.

message_file

Name of the message file to be checked.

when

Parameter that is set as `login` or `cwd=dir`. See the following example.

class

The `class` specification allows the message to be displayed only to members of a particular class.

Note – The `greeting` and `banner` keywords can also be used to send messages to users. See `ftppaccess(4)`.

Example—Configuring the README Option

```
readme README*      login
readme README*      cwd=*
```

The previous example states that any files that match `README*` are listed at login or when a directory is changed. Here is a sample login that is based on the settings that are used in that example.

```
% ftp earth
Connected to earth.
```

Composed September 5, 2002

```
220 earth FTP server ready.
Name (earth:rimmer): ftp
331 Guest login ok, send your complete e-mail address as password.
Password:
230-
230-Welcome to earth -- local time is Thu Jul 15 16:13:24
1999.
230-
230-You are number 1 out of a maximum of 10.
230-All transfers are logged.
230-
230-If your FTP client crashes or hangs shortly after login
please try
230-using a dash (-) as the first character of your
password. This will
230-turn off the informational messages that may be
confusing your FTP
230-client.
230-
230-Please send any comments to ftpadmin@earth.
230-
230 Guest login ok, access restrictions apply.
ftp> cd pub
250-Please read the file README
250- it was last modified on Thu Jul 15 16:12:25 1999 - 0
days ago
250 CWD command successful.
ftp> get README /tmp/README
200 PORT command successful.
150 Opening ASCII mode data connection for README (0
bytes).
226 ASCII Transfer complete.
ftp> quit
221 Goodbye.
```

Controlling Access to Files on the FTP Server

The FTP server access controls in this section supplement the standard file and directory access controls available with the Solaris 9 operating environment. Use the standard Solaris commands to restrict who can access, change, or upload files. See `chmod(1)`, `chown(1)`, and `chgrp(1)`.

▼ How to Control File Access Commands

To use the permission capabilities in `ftppaccess` to specify what type of user is allowed to perform which commands, do the following:

1. **Become superuser.**
2. **Add the following entries to the `ftppaccess` :**

command *yes|no* *typelist*

command

The commands `chmod`, `delete`, `overwrite`, `rename`, or `umask`

yes|no

Allows or disallows a user to issue a command

typelist

A comma-separated list of any of the keywords "anonymous," "guest," and "real"

Example—How to Control File Access Commands

The following are examples of permissions that are set for file access functions on FTP server.

```
chmod no anonymous, guest
delete no anonymous
overwrite no anonymous
rename no anonymous
umask no guest, anonymous
```

The preceding example states the following:

- Anonymous users are not allowed to delete, overwrite, or rename files.
- Guests and anonymous users are both prevented from changing access modes and resetting the `umask`.

Controlling Uploads and Downloads on the FTP Server

You can control uploads and downloads that are started to and from the FTP server by setting permissions on directories on the server. By default, uploads are not allowed for anonymous users. Be very careful when enabling anonymous uploads.

▼ How to Control Uploads to the FTP Server

Add the directives to the `ftppaccess` file to specify upload permissions and error messages for upload failures.

1. Become superuser.

2. Add the following entries to the `ftppaccess` file.

To enable users to upload files, add the following entry:

```
upload [absolute|relative] [class=<classname>]... [-] root-dir \  
dirglob yes|no owner group mode [dirs|nodirs] [<d_mode>]
```

```
path-filter typelist mesg allowed_charset {disallowed regexp...}
```

`upload`

Keyword that is applied to users who have a home directory (the argument to `chroot()`) of the *root-dir*. The *root-dir* can be specified as "*" to match any home directory.

`absolute|relative`

Parameter that specifies whether the *root-dir* directory paths are interpreted as absolute or relative to the current `chroot` directory.

`class`

Keyword that is used to specify any number of `class=<classname>` restrictions. If restrictions are specified, the upload clause only becomes effective if the current user is a member of one of the specified classes.

root-dir

User's root directory and the home directory for anonymous users.

dirglob

A pattern to match a directory name. An asterisk can be used in any place or alone to signify any directory.

`yes|no`

Variable that allows or disallows upload to the FTP server.

owner

Owner of files that are uploaded into `dirname`s.

group

Group that is associated with files that are uploaded into `dirname`s.

mode

Parameter that is used to specify access permissions for uploaded files. The default mode `0440` prevents the anonymous account from reading uploaded files.

<code>dirs nodirs</code>	Keyword that allows or disallows users to create subdirectories in a directory that is listed in <code>dirnames</code> .
<code>d_mode</code>	Optional mode that determines the permissions for a newly created directory.
<code>path-filter</code>	Keyword that controls the names of uploaded files.
<code>typelist</code>	A comma-separated list of any of the keywords "anonymous," "guest," and "real."
<code>msg</code>	Message file that is displayed fails to match the <code>regexp</code> criteria.
<code>allowed_charset {disallowed regexp...}</code>	Alphanumeric characters allowed or disallowed in file names.

Example—Controlling Uploads to the FTP Server

```
upload /export/home/ftp /incoming yes ftpadm ftpadmin 0440 nodirs
path-filter anonymous /etc/ftpd/filename.msg ^[-A-Za-z0-9._]*$ ^[.-]
```

The preceding example states the following:

- FTP user accounts that use `chroot` to `/export/home/ftp` can upload to the `/incoming` directory. Uploaded files are owned by user `ftpadm` and the group `ftpadm`. The mode is set to `0440` with the `nodirs` keyword to prevent anonymous users from creating subdirectories.
- For anonymous users, a file name is any sequence of A-Z, a-z, 0-9, . (dot), - (dash), or _ (underline). File names cannot start with a . (dot) or - (dash). If a file name fails this filter, the `/etc/ftpd/filename.msg` message is displayed if the FTP Administrator has created the message file. This message is followed by an FTP server error message.

Note – Ownership and permissions on a directory into which anonymous uploads are allowed should be tightly controlled. The FTP Administrator should be the owner of all files uploaded to the FTP server. You need to create an FTP Administrator when anonymous users are allowed to upload files. The directory should be owned by the user `ftpadm` and group `ftpadm` with permissions set to `3773`.

The access mode for files uploaded to the FTP server should be `0440`. The `0440` mode prevents the anonymous account from reading uploaded files. This restriction protects your server from becoming a staging area for third-party file distribution.

To make uploaded files available for distribution, the FTP Administrator can move files to a public directory.

▼ How to Control Downloads to the FTP Server

1. Become superuser.
2. Add the following entries to the `ftppaccess` file to prevent users from retrieving files.

```
noretrieve [absolute|relative] [class=classname]... [-] filename ...
```

`noretrieve`

Keyword that is used to deny retrieval of a particular file or files

`absolute|relative`

Parameter that specifies whether the *root-dir* directory paths are interpreted as absolute or relative to the current `chroot` directory

`class`

Keyword that is used to specify `class=<classname>` of users to which `noretrieve` restrictions apply

filename

Name of file the user is not permitted to retrieve

Example—Controlling Downloads to the FTP Server

```
noretrieve /etc/passwd
```

The preceding example states that all users are prevented from retrieving the `/etc/passwd` file.

Virtual Hosting

Virtual hosting allows the FTP server to support multiple domains on the same machine. Each virtual host requires a separate logical interface and IP address.

The FTP server supports two types of virtual hosting: *limited* and *complete*. With limited virtual hosting, the same configuration files are used for all virtual hosts. With complete virtual hosting, separate configuration files can be used for each virtual host.

Note – By default, real and guest users are not allowed to log in to virtual hosts. You can set the following `ftpaccess` directives to override the default.

To allow access to specific users:
`virtual address allow username`
To deny access to anonymous users:
`virtual address private username`

See `ftpaccess(4)` for further information.

▼ How to Enable Limited Virtual Hosting

Limited virtual hosting provides partial support for virtual FTP servers. You can enable support for limited virtual hosting by specifying the virtual root directory. If required, you can also set the following parameters for the virtual host in the `ftpaccess` file:

- `banner`
- `logfile`
- `email`
- `hostname`

All directives in the `ftpaccess` file are shared globally across all virtual servers.

1. **Become superuser.**
2. **Add the following entries to the `ftpaccess` file.**

```
virtual address root|banner|logfile path
virtual address hostname|email string
```

<code>virtual</code>	Keyword that is used to enable virtual server capabilities
<code>address</code>	IP address of the virtual server
<code>root</code>	The root directory of the virtual server
<code>banner</code>	Banner file that is displayed when a connection is made to the virtual server
<code>logfile</code>	Record of file transfers that are made to and from the virtual server
<code>path</code>	Variable that is used to specify the location of directories and files on the virtual server

<code>email</code>	Email address that is used in message files and in the <code>HELP</code> command
<code>hostname</code>	Name of the host that is shown in the greeting message or status command
<code>string</code>	Variable that is used to specify email or hostname parameters

Note – While it is possible to use `hostname` as the *address* of the virtual server, you are strongly encouraged to use the IPv4 address instead. DNS must be available when the FTP connection is received in order for `hostname` to be matched. For an IPv6 host, use the host name rather than the IPv6 address.

Example—Enabling Limited Virtual Hosting

```
virtual 10.1.2.3 root    /var/ftp/virtual/ftp-serv
virtual 10.1.2.3 banner  /var/ftp/virtual/ftp-serv/banner.msg
virtual 10.1.2.3 logfile /var/log/ftp/virtual/ftp-serv/xferlog
```

The preceding example sets the location of the root directory, banner, and logfile on a virtual FTP server.

Note – The `ftppaddhost(1M)` script with the `-l` option is provided to configure limited virtual hosts.

In the following example, `ftppaddhost` is run with `-l -b -x` options to configure limited virtual hosting with a test banner and the logfile `/var/ftp/virtual/10.1.2.3/xferlog` under a virtual root `/var/ftp/virtual/10.1.2.3`.

```
ftppaddhost -l -b -x /var/ftp/virtual/10.1.2.3/xferlog \
/var/ftp/virtual/10.1.2.3
```

▼ How to Enable Complete Virtual Hosting

Complete virtual hosting allows separate configuration files for each virtual domain. To enable complete support for virtual hosting on the FTP server, you can create or modify the following FTP configuration files for specific domains:

- `ftpaccess`
- `ftpusers`
- `ftpgroups`

Composed September 5, 2002

- `ftphosts`
- `ftpconversions`

For further information, see `ftpaccess(4)`, `ftpusers(4)`, `ftpgroups(4)`, `ftphosts(4)`, and `ftpconversions(4)`.

Note – If separate versions of the configuration files are unavailable, master versions of the files in the `/etc/ftpd` directory are used.

1. **Become superuser.**
2. **Add the following entry to the `/etc/ftpd/ftpservers` file.**

address */config-file-dir*

<i>address</i>	IP address of the virtual server
<i>config-file-dir</i>	Directory that contains the configuration files that are customized for the virtual host

Note – While it is possible to use `hostname` as the *address* of the virtual server, you are strongly encouraged to use the IPv4 address instead. DNS must be available when the FTP connection is received in order for `hostname` to be matched. For an IPv6 host, use the host name rather than the IPv6 address.

3. **To create a customized version of an FTP server configuration file for the virtual host, copy the master version of the file from `/etc/ftpd` to the `/config-file-dir` directory.**

For further information, see `ftpservers(4)`.

Example—Enabling Complete Virtual Hosting

```
#
# FTP Server virtual hosting configuration file
#

10.1.2.3 /net/inet/virtual/somedomain/
10.1.2.4 /net/inet/virtual/anotherdomain/
```

The preceding example specifies the IP addresses for two different domains on the virtual server.

Note – The `ftppaddhost(1M)` script with the `-c` option is provided to configure complete virtual hosts.

In the following example, `ftppaddhost` is run with `-l -b -x` options to configure limited virtual hosting with a test banner and the logfile `/var/ftp/virtual/10.1.2.3/xferlog` under a virtual root `/var/ftp/virtual/10.1.2.3`.

```
ftppaddhost -l -b -x /var/ftp/virtual/10.1.2.3/xferlog \  
/var/ftp/virtual/10.1.2.3
```

Starting the FTP Server Automatically

The FTP server can be started in one of two ways:

- As a `nowait` that is server started from the `inetd.conf` file
- As a standalone server that is started from the command line or by a startup script

Starting an FTP Server From `inetd.conf`

You can add a `nowait` entry in `inetd.conf` file to start the FTP server. If the site handles many connections, the FTP daemon can also be run in standalone mode. For more information, see `inetd.conf(4)`. See also `in.ftpd(1M)` for information on additional command-line options.

▼ How to Start an FTP Server From `inetd.conf`

1. **Become superuser.**
2. **Add a `nowait` entry to the `inetd.conf` file:**

```
ftp  stream  tcp6  nowait  root  /usr/sbin/in.ftpd  in.ftpd -a
```

Note – The `-a` option is specified to enable the use of the `ftpaccess` file.

3. **Signal `inetd` to reread the `inetd.conf` file.**

```
# pkill -HUP inetd
```

Starting a Standalone FTP Server

The FTP server can also be run independently of the `inetd.conf` as a *standalone* server.

A standalone server always has the quickest possible response time, and is intended for large servers that are dedicated to providing FTP service. The standalone server provides low connection latency for dedicated servers because the standalone system never has to be restarted. The standalone server is always running—even during off-peak hours—waiting indefinitely for connections.

▼ How to Start a Standalone FTP Server

1. **Become superuser.**
2. **Add a # sign at the start of the ftp service line in the `inetd.conf` file to comment out the entry.**
3. **Signal `inetd` to reread the `inetd.conf` file.**

```
# pkill -HUP inetd
```

4. **Start the standalone FTP server.**

```
# /usr/sbin/in.ftpd -a -S
```

Add the line to an FTP server startup script. See “Run Control Scripts” in *System Administration Guide, Volume 1* for information on creating a system startup script.

Shutting Down the FTP Server

The `ftpsht(1M)` command closes down the FTP server at a particular time.

When you run `ftpsht`, a file is generated from command-line options that specify when shutdown occurs, the point at which new connections are refused, and when existing connections are dropped. Users are notified of a server shutdown based on this information. The location of the file that is created by `ftpsht` is specified by the shutdown directive in the `ftpaccess` file.

▼ How to Shut Down the FTP Server

Follow the steps in this procedure to run `ftpsht` and to add the shutdown directive to the `ftpaccess` file.

1. Become superuser.

2. Add the following entries to the `ftpaccess` file.

`shutdown path`

`shutdown`

Keyword that is used to specify the *path* to a file that is checked regularly for whether the FTP server is scheduled to be shut down

path

Location of the file that was created by `ftpshut` command

3. Run the `ftpshut` command.

`ftpshut [ -V ] [ -l min] [ -d min] time [warning-message...]`

`ftpshut`

Command that provides a procedure for notifying users that the FTP server is shutting down.

`-V`

Option that is specified to display copyright and version information, then terminate

`-l`

Flag that is used to adjust the time that new connections to the FTP server are denied

`-d`

Flag that is used to adjust the time that existing connections to the FTP server are disconnected.

`time`

Shutdown time that is specified by the word `now` for immediate shutdown, or in one of two formats (+ *number* or *HHMM*) for a future shutdown

`[warning-message...]`

Shutdown notification message

4. Use the `ftprestart` command to restart the FTP server after shutdown.

For further information, see `ftpshut(1M)`, `ftpaccess(4)`, and `ftprestart(1M)`.

Debugging the FTP Server

This section describes some of the ways to debug problems with the FTP server.

▼ How to Check syslogd for FTP Server Messages

The FTP server writes messages that are useful for debugging to the location that is specified for daemon messages in the `/etc/syslog.conf` file. If a problem occurs with the FTP server, check this file first for such messages.

The FTP server messages are controlled by facility daemon and level information. To send messages from the FTP server to `/var/adm/message` and have `syslogd` reread its configuration file, follow these instructions:

1. **Add an entry such as the following to the `/etc/syslog.conf` file.**

```
daemon.info /var/adm/message
```

2. **Signal `syslogd` to reread its configuration file.**

```
# pkill -HUP syslogd
```

This action causes informational messages from the FTP Server to be written to `/var/adm/messages`.

▼ How to Use greeting text to Verify `ftppaccess`

To use the greeting text capability to check that the correct `ftppaccess` file is being used, do the following:

1. **Add the following directive to the `ftppaccess` file.**

```
greeting text message
```

2. **Connect to the FTP server.**

3. **If the message fails to appear, do the following:**

- a. **Confirm that the `ftppaccess` file is in the correct location. Use the `strings(1)` command to obtain the location of the file from the FTP server binary.**

```
# strings /usr/sbin/in.ftpd | grep "^/*.*ftppaccess"
```

- b. **Check the `ftpservers` file to see if virtual hosting has been configured.**

For further information, see `ftppaccess(4)`, `ftpservers(4)`, `strings(1)`, `syslog.conf(4)`, and `pkill(1)`.

▼ How to Check the Commands Executed by FTP Users

To see what commands are being executed by FTP users, use the log commands logging capability in `ftppaccess`.

Composed September 5, 2002

1. Add the following directive to the `ftppaccess` file to log individual commands by users who are specified in *typelist*.

```
log commands typelist
```

2. Check messages that are written to the location specified in the `/etc/syslog.conf`.

Accessing Remote Systems (Tasks)

This chapter describes all the tasks that are required to log in to remote systems and work with their files. This is a list of the step-by-step instructions in this chapter.

- “How to Search for and Remove `.rhosts` Files” on page 662
- “How to Find Out If a Remote System Is Operating” on page 663
- “How to Find Who Is Logged In to a Remote System” on page 664
- “How to Log In to a Remote System (`rlogin`)” on page 665
- “How to Log Out From a Remote System (`exit`)” on page 665
- “How to Open an `ftp` Connection to a Remote System” on page 667
- “How to Close an `ftp` Connection to a Remote System” on page 668
- “How to Copy Files From a Remote System (`ftp`)” on page 668
- “How to Copy Files to a Remote System (`ftp`)” on page 670
- “How to Copy Files Between a Local and a Remote System (`rcp`)” on page 675

This chapter provides tasks that are described in the following table to log in and copy files from remote systems.

TABLE 44-1 Task Map: Accessing Remote Systems

Task	Description	For Instructions
Log in to a remote system (<code>rlogin</code>)	■ Remove <code>.rhosts</code> files. ■ Use the <code>rlogin</code> command to access a remote system.	“How to Search for and Remove <code>.rhosts</code> Files” on page 662 “How to Find Out If a Remote System Is Operating” on page 663 “How to Find Who Is Logged In to a Remote System” on page 664 “How to Log In to a Remote System (<code>rlogin</code>)” on page 665 “How to Log Out From a Remote System (<code>exit</code>)” on page 665

TABLE 44-1 Task Map: Accessing Remote Systems (Continued)

Task	Description	For Instructions
Log in to a remote system (<code>ftp</code>)	■ Open and close an <code>ftp</code> connection. ■ Copy files to and from a remote system.	"How to Open an <code>ftp</code> Connection to a Remote System" on page 667 "How to Close an <code>ftp</code> Connection to a Remote System" on page 668 "How to Copy Files From a Remote System (<code>ftp</code>)" on page 668 "How to Copy Files to a Remote System (<code>ftp</code>)" on page 670
Copy remote files with <code>rcp</code>	Use the <code>rcp</code> command to copy files to and from a remote system.	"How to Copy Files Between a Local and a Remote System (<code>rcp</code>)" on page 675

Logging In to a Remote System (`rlogin`)

The `rlogin` command enables you to log in to a remote system. After you are logged in, you can navigate through the remote file system and manipulate its contents (subject to authorization), copy files, or execute remote commands.

If the system you are logging in to is in a remote domain, be sure to append the domain name to the system name. In this example, `SOLAR` is the name of the remote domain:

```
rlogin pluto.SOLAR
```

Also, you can interrupt a remote login operation at any time by typing `Control-d`.

Authentication for Remote Logins (`rlogin`)

Authentication (establishing who you are) for `rlogin` operations can be performed either by the remote system or by the network environment.

The main difference between these forms of authentication lies in the type of interaction they require from you and the way they are established. If a remote system tries to authenticate you, you are prompted for a password, unless you set up the `/etc/hosts.equiv` or `.rhosts` file. If the network tries to authenticate you, you are not asked for a password, because the network already knows who you are.

When the remote system attempts to authenticate you, it relies on information in its local files, specifically if one of the following is true:

- Your system name and user name appear in the remote system's `/etc/hosts.equiv` file.
- Or:
- Your system name and user name appear in the remote user's `.rhosts` file, under the remote user's home directory.

Network authentication relies on one of these two methods:

- A "trusting network environment" that has been set up with your local network information service and the automounter
- One of the network information services that is pointed to by the remote system's `/etc/nsswitch.conf` file contains information about you

Note – Network authentication generally supersedes system authentication.

`/etc/hosts.equiv` File

The `/etc/hosts.equiv` file contains a list of trusted hosts for a remote system, one per line. If a user attempts to log in remotely (using `rlogin`) from one of the hosts that is listed in this file, and if the remote system can access the user's password entry, the remote system allows the user to log in without a password.

A typical `hosts.equiv` file has the following structure:

```
host1
host2 user_a
+@group1
-@group2
```

When a simple entry for a host is made in `hosts.equiv`, such as the previous entry for `host1`, it means that the host is trusted, and so is any user at that machine.

If the user name is also mentioned, as in the second entry in the example, then the host is trusted only if the specified user is attempting access.

A group name that is preceded by a plus sign (+) means that all the machines in that netgroup are considered trusted.

A group name that is preceded by a minus sign (–) means that none of the machines in that netgroup is considered trusted.

Security Risks When Using the /etc/hosts.equiv File

The `/etc/hosts.equiv` file presents a security risk. If you maintain a `/etc/hosts.equiv` file on your system, you should include only trusted hosts in your network. The file should not include any host that belongs to a different network, or any machines that are in public areas. For example, do not include a host that is located in a terminal room.

The use of hosts that are not trusted can create a serious security problem. Either replace the `/etc/hosts.equiv` file with a correctly configured one, or remove the file altogether.

A single line of `+` in the `/etc/hosts.equiv` file indicates that every known host is trusted.

.rhosts File

The `.rhosts` file is the user equivalent of the `/etc/hosts.equiv` file. This file contains a list of host-user combinations, rather than hosts in general. If a host-user combination is listed in this file, the specified user is granted permission to log in remotely from the specified host without having to supply a password.

Note that a `.rhosts` file must reside at the top level of a user's home directory. `.rhost` files that are located in subdirectories are not consulted.

Users can create `.rhosts` files in their home directories. Using the `.rhosts` file is another way to allow trusted access between users' own accounts on different systems without using the `/etc/hosts.equiv` file.

Security Risks When Using the .rhosts File

Unfortunately, the `.rhosts` file presents a major security problem. While the `/etc/hosts.equiv` file is under the system administrator's control and can be managed effectively, any user can create a `.rhosts` file that grants access to whomever the user chooses without the system administrator's knowledge.

In a situation in which all of the users' home directories are on a single server and only certain people have superuser access on that server, a good way to prevent a user from using a `.rhosts` file is to create an empty file as superuser in their home directory. You would then change the permissions in this file to `000` so that it would be difficult to change it, even as superuser. This change would effectively prevent a user from risking system security by using a `.rhosts` file irresponsibly. The change would not, however, solve anything if the user is able to change the effective path to his or her home directory.

The only secure way to manage `.rhosts` files is to completely disallow them. See “How to Search for and Remove `.rhosts` Files” on page 662 for detailed instructions. As system administrator, you can check the system often for violations of this policy. One possible exception to this policy is for the root account—you might need to have a `.rhosts` file to perform network backups and other remote services.

Linking Remote Logins

If your system is configured properly, you can link remote logins. For example, a user on `earth` logs in to `jupiter`, and from there decides to log in to `pluto`.

The user could have logged out of `jupiter` and then logged in directly to `pluto`, but this type of linking can be more convenient.

To link remote logins without having to supply a password, you must have the `/etc/hosts.equiv` or `.rhosts` file set up correctly.

Direct or Indirect Remote Logins

The `rlogin` command allows you to log in to a remote system directly or indirectly.

A direct remote login is attempted with the default user name, that is, the user name of the individual who is currently logged in to the local system. This is the most common form of remote login.

An indirect remote login is attempted with a different user name, which is supplied during the remote login operation. This is the type of remote login you might attempt from a workstation that you borrowed temporarily. For instance, if you were in a coworker’s office and needed to examine files in your home directory, you might log in to your system remotely, from your coworker’s system. However, you would perform an indirect remote login, supplying your own user name.

The dependencies between direct and indirect logins and authentication methods are summarized in the following table.

TABLE 44–2 Dependencies Between Login Method and Authentication Method (`rlogin`)

Type of Login	User Name Supplied By	Authentication	Password
Direct	System	Network	None
		System	Required
Indirect	User	Network	None
		System	Required

What Happens After You Log In Remotely

When you log in to a remote system, the `rlogin` command attempts to find your home directory. If the `rlogin` command can't find your home directory, it assigns you to the remote system's root (`/`) directory. For example:

```
Unable to find home directory, logging in with /
```

However, if the `rlogin` command finds your home directory, it sources both your `.cshrc` and `.login` files. Therefore, after a remote login, your prompt is your standard login prompt, and the current directory is the same as when you log in locally.

For example, if your usual prompt displays your system name and working directory, and when you log in, your working directory is your home directory, your login prompt resembles the following:

```
earth(/home/smith):
```

Then when you log in to a remote system, you see a similar prompt and your working directory is your home directory, regardless of the directory from which you entered the `rlogin` command:

```
earth(/home/smith): rlogin pluto
.
.
.
pluto(/home/smith):
```

The only difference is that the name of the remote system would substitute for your local system at the beginning of the prompt. The remote file system is parallel to your home directory.

Effectively, if you change directory to `/home` and then run `ls`, you see the following:

```
earth(home/smith): cd ..
earth(/home): ls
smith jones
```

▼ How to Search for and Remove `.rhosts` Files

1. Become superuser.
2. Search for and remove `.rhosts` files by using the `find(1)` command.

```
# find home-directories -name .rhosts -print -exec rm{}
```

<i>home-directories</i>	Identifies the path to a directory where users' home directories are located. Note that you can enter multiple paths to search more than one home directory at a time.
<code>-name .rhosts</code>	Identifies the file name.
<code>-print</code>	Prints the current path name.
<code>-exec rm {} \;</code>	Tells the <code>find</code> command to apply the <code>rm</code> command to all files that are identified by using the matching file name.

The `find` command starts at the designated directory and searches for any file that is named `.rhosts`. If it finds such as file, `find` prints the path on the screen and removes it.

Example—Searching for and Removing `.rhosts` Files

The following example searches and removes `.rhosts` files in all the user's home directories that are located in the `/export/home` directory.

```
# find /export/home -name .rhosts -print | xargs -i -t rm{}
```

▼ How to Find Out If a Remote System Is Operating

Find out if a remote system is operating by using the `ping` command.

```
$ ping system-name | ip-address
```

<i>system-name</i>	The name of the remote system
<i>ip-address</i>	The IP address of the remote system

The `ping` command returns one of three messages:

Status Message	Explanation
<i>system-name</i> is alive	The system can be accessed over the network.
ping:unknown host <i>system-name</i>	The system name is unknown.
ping:no answer from <i>system-name</i>	The system is known, but is not currently operating.

If the system you “ping” is located in a different domain, the return message can also contain routing information, which you can ignore.

The `ping` command has a timeout of 20 seconds. Effectively, if it does not receive a response within 20 seconds, it returns the third message. You can force `ping` to wait longer (or less) by typing a *time-out* value, in seconds:

```
$ ping system-name | ip-address time-out
```

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

▼ How to Find Who Is Logged In to a Remote System

Find who is logged in to a remote system by using the `rusers(1)` command.

```
$ rusers [-l] remote-system-name
```

<code>rusers</code>	(No options) Displays the name of the system, followed by the name of users who are currently logged in to it, including root
<code>-l</code>	Displays additional information about each user: the user's login window, login time and date, amount of time logged in, and the name of the remote system from which the user logged on

Example—Finding Who Is Logged In to a Remote System

The following example shows the short output of `rusers`.

```
$ rusers pluto
pluto    smith  jones
```

In the following example, the long version of `rusers` shows that two users are logged in to the remote system `starbug`. The first user logged in from the system console on September 10 and has been logged on for 137 hours and 15 minutes. The second user logged in from a remote system, `mars`, on September 14.

```
$ rusers -l starbug
root      starbug:console      Sep 10 16:13  137:15
rimmer    starbug:pts/0        Sep 14 14:37      (mars)
```


▼ How to Log In to a Remote System (`rlogin`)

Log in to a remote system by using the `rlogin(1)` command.

```
$ rlogin [-l user-name] system-name
```

<code>rlogin</code>	(No options) Logs you in to the remote system <i>directly</i> , effectively, with your current user name
<code>-l <i>user-name</i></code>	Logs you into the remote system <i>indirectly</i> , effectively, with the user name you supply

If the network attempts to authenticate you, you are not prompted for a password. If the remote system attempts to authenticate you, you are asked to provide a password.

If the operation succeeds, the `rlogin` command displays brief information about your latest remote login to that system, the version of the operating system that is running on the remote system, and whether you have mail waiting for you in your home directory.

Example—Logging In to a Remote System (`rlogin`)

The following example shows the output of a direct remote login to `pluto`. The user has been authenticated by the network.

```
$ rlogin starbug
Last login: Mon Jul 12 09:28:39 from venus
Sun Microsystems Inc.   SunOS 5.8           February 2000
starbug:
```

The following example shows the output of an indirect remote login to `pluto`, with the user being authenticated by the remote system.

```
$ rlogin -l smith pluto
password: user-password
Last login: Mon Jul 12 11:51:58 from venus
Sun Microsystems Inc.   SunOS 5.8           February 2000
starbug:
```

▼ How to Log Out From a Remote System (`exit`)

Log out from a remote system by using the `exit(1)` command.

```
$ exit
```

Example—Logging Out From a Remote System (`exit`)

This example shows the user `smith` logging out from the system `pluto`.

```
$ exit
pluto% logout
Connection closed.
earth%
```

Logging In to a Remote System (`ftp`)

The `ftp` command opens the user interface to the Internet's File Transfer Protocol. This user interface, called the command interpreter, enables you to log in to a remote system and perform a variety of operations with its file system. The principal operations are summarized in the following table.

The main benefit of `ftp` over `rlogin` and `rcp` is that `ftp` does not require the remote system to be running UNIX. The remote system does, however, need to be configured for TCP/IP communications. However, `rlogin` provides access to a richer set of file manipulation commands than `ftp` provides.

Authentication for Remote Logins (`ftp`)

Authentication for `ftp` remote login operations can be established by one of the following methods:

- Including your password entry in the remote system's `/etc/passwd` file or equivalent network information service map or table
- Establishing an anonymous `ftp` account on the remote system

Essential `ftp` Commands

TABLE 44-3 Essential `ftp` Commands

Command	Description
<code>ftp</code>	Accesses the <code>ftp</code> command interpreter.
<code>ftp remote-system</code>	Establishes an <code>ftp</code> connection to a remote system. For instructions, see "How to Open an <code>ftp</code> Connection to a Remote System" on page 667.

TABLE 44-3 Essential `ftp` Commands (Continued)

Command	Description
<code>open</code>	Logs in to the remote system from the command interpreter.
<code>close</code>	Logs out of the remote system and returns to the command interpreter.
<code>bye</code>	Quits the <code>ftp</code> command interpreter.
<code>help</code>	Lists all <code>ftp</code> commands or, if a command name is supplied, briefly describes what the command does.
<code>reset</code>	Re-synchronizes the command-reply sequencing with the remote <code>ftp</code> server.
<code>ls</code>	Lists the contents of the remote working directory.
<code>pwd</code>	Displays the name of the remote working directory.
<code>cd</code>	Changes the remote working directory.
<code>lcd</code>	Changes the local working directory.
<code>mkdir</code>	Creates a directory on the remote system.
<code>rmdir</code>	Deletes a directory on the remote system.
<code>get</code> , <code>mget</code>	Copies a file (or multiple files) from the remote working directory to the local working directory.
<code>put</code> , <code>mput</code>	Copies a file (or multiple files) from the local working directory to the remote working directory.
<code>delete</code> , <code>mdelete</code>	Deletes a file (or multiple files) from the remote working directory.

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

▼ How to Open an `ftp` Connection to a Remote System

1. Ensure that you have `ftp` authentication.

You must have `ftp` authentication, as described in “Authentication for Remote Logins (`ftp`)” on page 666.

2. Open a connection to a remote system by using the `ftp` command.

```
$ ftp remote-system
```

If the connection succeeds, a confirmation message and prompt are displayed.

3. Type your user name.

```
Name (remote-system:user-name) : user-name
```

4. If prompted, type your password.

```
331 Password required for user-name:  
Password: password
```

If the system you are accessing has an established anonymous ftp account, you are prompted for an email address for the password. If the ftp interface accepts your password, it displays a confirmation message and the (ftp>) prompt.

You can now use any of the commands that are supplied by the ftp interface, including help. The principal commands are summarized in Table 44–3.

Example—Opening an ftp Connection to a Remote System

This ftp session was established by the user smith on the remote system pluto:

```
$ ftp pluto  
Connected to pluto.  
220 pluto FTP server ready.  
Name (pluto:smith): smith  
331 Password required for smith:  
Password: password  
230 User smith logged in.  
ftp>
```

▼ How to Close an ftp Connection to a Remote System

Close an ftp connection to a remote system by using the bye command.

```
ftp> bye  
221-You have transferred 0 bytes in 0 files.  
221-Total traffic for this sessions was 172 bytes in 0 transfers.  
221-Thanks you for using the FTP service on spdev.  
221 Goodbye.
```

A goodbye message appears, followed by your usual shell prompt.

▼ How to Copy Files From a Remote System (ftp)

1. Change to a directory on the local system where you want the files from the remote system to be copied.

```
$ cd target-directory
```

2. Establish an ftp connection.

See “How to Open an ftp Connection to a Remote System” on page 667.

3. Change to the source directory.

```
ftp> cd source-directory
```

If your system is using the automounter, the home directory of the remote system’s user appears parallel to yours, under /home.

4. Ensure that you have read permission for the source files.

```
ftp> ls -l
```

5. Set the transfer type to binary.

```
ftp> binary
```

6. To copy a single file, use the get command.

```
ftp> get filename
```

7. To copy multiple files at once, use the mget command.

```
ftp> mget filename [filename ...]
```

You can supply a series of individual file names and you can use wildcard characters. The mget command copies each file individually, asking you for confirmation each time.

8. Close the ftp connections.

```
ftp> bye
```

Examples—Copying Files From a Remote System (ftp)

In this example, the user kryten opens an ftp connection to the system pluto, and uses the get command to copy a single file from the /tmp directory.

```
$ cd $HOME
ftp pluto
Connected to pluto.
220 pluto FTP server (SunOS 5.8) ready.
Name (pluto:kryten): kryten
331 Password required for kryten.
Password: xxx
230 User kryten logged in.
ftp> cd /tmp
250 CWD command successful.
ftp> ls
200 PORT command successful.
150 ASCII data connection for /bin/ls (129.152.221.238,34344)
(0 bytes).
dtddbcache_:0
```

```
filea
files
ps_data
speckeyds.lock
226 ASCII Transfer complete.
53 bytes received in 0.022 seconds (2.39 Kbytes/s)
ftp> get filea
200 PORT command successful.
150 ASCII data connection for filea (129.152.221.238,34331)
(0 bytes).
221 Goodbye.
```

In this example, the same user kryten uses the `mget` command to copy a set of files from the `/tmp` directory to his home directory. Note that kryten can accept or reject individual files in the set.

```
$ ftp> cd /tmp
250 CWD command successful.
ftp> ls files
200 PORT command successful.
150 ASCII data connection for /bin/ls (129.152.221.238,34345)
(0 bytes).
fileb
filec
filed
remote: files
21 bytes received in 0.015 seconds (1.36 Kbytes/s)
ftp> cd files
250 CWD command successful.
ftp> mget file*
mget fileb? y
200 PORT command successful.
150 ASCII data connection for fileb (129.152.221.238,34347)
(0 bytes).
226 ASCII Transfer complete.
mget filec? y
200 PORT command successful.
150 ASCII data connection for filec (129.152.221.238,34348)
(0 bytes).
226 ASCII Transfer complete.
mget filed? y
200 PORT command successful.
150 ASCII data connection for filed (129.152.221.238,34351)
(0 bytes).
226 ASCII Transfer complete.200 PORT command successful.
ftp> bye
221 Goodbye.
```

▼ How to Copy Files to a Remote System (ftp)

1. Change to the source directory on the local system.

The directory from which you type the `ftp` command is the local working directory, and thus the source directory for this operation.

2. Establish an ftp connection.

See “How to Open an ftp Connection to a Remote System” on page 667.

3. Change to the target directory.

```
ftp> cd target-directory
```

Remember, if your system is using the automounter, the home directory of the remote system’s user appears parallel to yours, under /home.

4. Ensure that you have write permission to the target directory.

```
ftp> ls -l target-directory
```

5. Set the transfer type to binary.

```
ftp> binary
```

6. To copy a single file, use the put command.

```
ftp> put filename
```

7. To copy multiple files at once, use the mput command.

```
ftp> mput filename [filename ...]
```

You can supply a series of individual file names and you can use wildcard characters. The mput command copies each file individually, asking you for confirmation each time.

8. To close the ftp connection, type bye.

```
ftp> bye
```

Examples—Copying Files to a Remote System (ftp)

In this example, the user kryten opens an ftp connection to the system pluto, and uses the put command to copy a file from his or her system to the /tmp directory on system pluto.

```
$ cd /tmp
ftp pluto
Connected to pluto.
220 pluto FTP server (SunOS 5.8) ready.
Name (pluto:kryten): kryten
331 Password required for kryten.
Password: xxx
230 User kryten logged in.
ftp> cd /tmp
250 CWD command successful.
ftp> put filef
200 PORT command successful.
150 ASCII data connection for filef (129.152.221.238,34356).
226 Transfer complete.
```

Composed September 5, 2002

```
ftp> ls
200 PORT command successful.
150 ASCII data connection for /bin/ls (129.152.221.238,34357) (0 bytes).
dtddbcache_:0
filea
filef
files
ps_data
speckeyd.lock
226 ASCII Transfer complete.
60 bytes received in 0.058 seconds (1.01 Kbytes/s)
ftp> bye
221 Goodbye.
```

In this example, the same user kryten uses the mput command to copy a set of files from his or her home directory to pluto's /tmp directory. Note that kryten can accept or reject individual files in the set.

```
$ cd $HOME/testdir
$ ls
test1  test2  test3
$ ftp pluto
Connected to pluto.
220 pluto FTP server (SunOS 5.8) ready.
Name (pluto:kryten): kryten
331 Password required for kryten.
Password: xxx
230 User kryten logged in.
ftp> cd /tmp
250 CWD command successful.
ftp> mput test*
mput test1? y
200 PORT command successful.
150 ASCII data connection for test1 (129.152.221.238,34365).
226 Transfer complete.
mput test2? y
200 PORT command successful.
150 ASCII data connection for test2 (129.152.221.238,34366).
226 Transfer complete.
mput test3? y
200 PORT command successful.
150 ASCII data connection for filef (129.152.221.238,34356).
226 Transfer complete.
ftp> bye
221 Goodbye.
```

Remote Copying With `rcp`

The `rcp` command copies files or directories between a local and a remote system or between two remote systems. You can use this command from a remote system (after logging in with the `rlogin` command) or from the local system (without logging in to a remote system).

With `rcp`, you can perform the following remote copy operations:

- Copy a file or directory from your system to a remote system
- Copy a file or directory from a remote system to your local system
- Copy a file or directory between remote systems from your local system

If you have the automounter running, you can perform these remote operations with the `cp` command. However, the range of `cp` is constrained to the virtual file system that is created by the automounter and to operations relative to a user's home directory. Because `rcp` performs the same operations without these constraints, this section describes only the `rcp` versions of these tasks.

Security Considerations for Copy Operations

To copy files or directories between systems, you must have permission to log in and copy files.



Caution – Both the `cp` and `rcp` commands can overwrite files without warning. Ensure that file names are correct before executing the command.

Specifying Source and Target

With the `rcp` command in the C shell, you can specify source (the file or directory you want to copy) and target (the location into which you will copy the file or directory) with either absolute or abbreviated path names.

	Absolute Path Names	Abbreviated Path Names
From Local System	<code>mars:/home/jones/myfile.txt</code>	<code>~jones/myfile.txt</code>

	Absolute Path Names	Abbreviated Path Names
After Remote Login	/home/jones/myfile.txt	~jones/myfile.txt

Absolute path names identify files or directories that are mounted on a particular system. In the previous example, the first absolute path name identifies a file (`MyFile.txt`) on the `mars` system. Abbreviated path names identify files or directories relative to a user's home directory, wherever it might reside. In the previous first example, the abbreviated path name identifies the same file, `MyFile.txt`, but uses "`~`" symbol to indicate the `jones` home directory:

`~ = mars:/home/jones`

The examples on the second line demonstrate the user of absolute and abbreviated path names after a remote login. No difference is evident for the abbreviated path name. However, because the remote login operation mounted the `jones` home directory onto the local system (parallel to the local user's home directory), the absolute path name no longer requires the system name `mars`. For more information about how a remote login operation mounts another user's home directory, see "What Happens After You Log In Remotely" on page 662.

The following table provides a sample of absolute and abbreviated path names that are recognized by the C shell. The sample uses the following terminology:

- Working directory — The directory from which the `rcp` command is entered. Can be remote or local.
- Current user — The user name under which the `rcp` command is entered.

TABLE 44-4 Allowed Syntaxes for Directory and File Names

Logged in to	Syntax	Description
Local system	<code>.</code>	The local working directory
	<code>path/filename</code>	The <i>path</i> and <i>filename</i> in the local working directory
	<code>~</code>	The current user's home directory
	<code>~/path/filename</code>	The <i>path</i> and <i>filename</i> beneath the current user's home directory
	<code>~user</code>	The home directory of <i>user</i>
	<code>~user/path/filename</code>	The <i>path</i> and <i>filename</i> beneath the home directory of <i>user</i>
	<code>remote-system:path/filename</code>	The <i>path</i> and <i>filename</i> in the remote working directory

TABLE 44-4 Allowed Syntaxes for Directory and File Names (Continued)

Logged in to	Syntax	Description
Remote system	.	The remote working directory
	<i>filename</i>	The <i>filename</i> in the remote working directory
	<i>path/filename</i>	The <i>path</i> and <i>filename</i> in the remote working directory
	~	The current user's home directory
	~/ <i>path/filename</i>	The <i>path</i> and <i>filename</i> in the current user's home directory
	~ <i>user</i>	The home directory of <i>user</i>
	~/ <i>user/path/filename</i>	The <i>path</i> and <i>filename</i> beneath the home directory of <i>user</i>
	<i>local-system:path/filename</i>	The <i>path</i> and <i>filename</i> in the local working directory

▼ How to Copy Files Between a Local and a Remote System (r_{cp})

1. Ensure that you have permission to copy.

You should at least have read permission on the source system and write permission on the target system.

2. Determine the location of the source and target.

If you don't know the path of the source or target, you can first log in to the remote system with the `rlogin` command, as described in "How to Log In to a Remote System (rlogin)" on page 665. Then, navigate through the remote system until you find the location. You can then perform the next step without logging out.

3. Copy the file or directory.

```
$ rcp [-r] source-file\directory target-file\directory
```

`rcp` (No options) Copies a single file from the source to the target.

`-r` Copies a directory from the source to the target.

This syntax applies whether you are logged in to the remote system or in to the local system. Only the path name of the file or directory changes, as described in Table 44-4 and as illustrated in the following examples.

You can use the “~” and “.” characters to specify the path portions of the local file or directory names. Note, however, that “~” applies to the current user, not the remote system, and that “.” applies to system you are logged in to. For explanations of these symbols, see Table 44–4.

Examples—Copying Files Between a Local and a Remote System (rcp)

Here are several examples of using rcp to copy files to and from local and remote systems.

Using rcp to Copy a Remote File to a Local System

In this example, rcp is used to copy the file `letter.doc` from the `/home/jones` directory of the remote system `pluto` to the working directory (`/home/smith`) on the local system, `earth`:

```
earth(/home/smith): rcp pluto:/home/jones/letter.doc .
```

In this instance, the rcp operation is performed without a remote login. Here, the “.” symbol at the end of the command line refers to the local system, not the remote system.

The target directory is the also local user’s home directory, so it can also be specified with the “~” symbol.

```
earth(home/smith): rcp pluto:/home/jones/letter.doc ~
```

Using rlogin and rcp to Copy a Remote File to a Local System

In this example, the rcp operation is run after the rlogin command is executed to copy a file from a remote to a local system. Although the flow of the operation is the same as that of the previous example, the paths change to allow for the remote login:

```
earth(/home/smith): rlogin pluto
.
.
.
pluto(/home/jones): rcp letter.doc ~
```

Using the “.” symbol at the end of the command line would be inappropriate in this instance. Because of the remote login, the symbol would simply refer to the remote system—essentially directing rcp to create a duplicate file. The “~” symbol, however, refers to the current user’s home directory, even when the login is to a remote system.

Using rcp to Copy a Local File to a Remote System

In this example, `rcp` is used to copy the file `notice.doc` from the home directory (`/home/smith`) of the local system `earth` to the `/home/jones` directory of the remote system, `pluto`:

```
earth(/home/smith): rcp notice.doc pluto:/home/jones
```

Because no remote file name is provided, the file `notice.doc` is copied into the `/home/jones` directory with the same name.

In this instance, the `rcp` operation from the previous example is repeated, but `rcp` is entered from a different working directory on the local system (`/tmp`). Note the use of the `"~"` symbol to refer to the current user's home directory:

```
earth(/tmp): rcp ~/notice.doc pluto:/home/jones
```

Using rlogin and rcp to Copy a Local File to a Remote System

In this example, the `rcp` operation is run after the `rlogin` command is executed to copy a local file to a remote directory. Although the flow of the operation is the same as that of the previous example, the paths change to allow for the remote login.

```
earth(/home/smith): rlogin pluto
.
.
.
pluto(/home/jones): rcp ~/notice.doc .
```

In this instance, the `"~"` symbol can be used to denote the current user's home directory, even though it is on the local system. The `"."` symbol refers to the working directory on the remote system because the user is logged in to the remote system. Here is an alternative syntax that performs the same operation:

```
pluto(/home/jones): rcp earth:/home/smith/notice.doc /home/jones
```

Composed September 5, 2002

CHAPTER **45**

Monitoring Network Services Topics

Chapter 46

Provides step-by-step instructions for monitoring network services

Composed September 5, 2002

Monitoring Network Performance (Tasks)

This chapter describes the how to monitor network performance. This is a list of the step-by-step instructions in this chapter.

- “How to Check the Response of Hosts on the Network” on page 682
- “How to Send Packets to Hosts on the Network” on page 683
- “How to Capture Packets From the Network” on page 683
- “How to Check the Network Status” on page 683
- “How to Display NFS Server and Client Statistics” on page 686

Monitoring Network Performance

Table 46-1 describes the commands available for monitoring network performance.

TABLE 46-1 Network Monitoring Commands

Command	Description
ping	Look at the response of hosts on the network.
spray	Test the reliability of your packet sizes. This command can tell you whether packets are being delayed or dropped.
snoop	Capture packets from the network and trace the calls from each client to each server.
netstat	Display network status, including state of the interfaces that are used for TCP/IP traffic, the IP routing table, and the per-protocol statistics for UDP, TCP, ICMP, and IGMP.
nfsstat	Display a summary of server and client statistics that can be used to identify NFS problems.

▼ How to Check the Response of Hosts on the Network

Check the response of hosts on the network with the `ping` command.

```
$ ping hostname
```

If you suspect a physical problem, you can use `ping` to find the response time of several hosts on the network. If the response from one host is not what you would expect, you can investigate that host. Physical problems could be caused by the following:

- Loose cables or connectors
- Improper grounding
- Missing termination
- Signal reflection

For more information about this command, see `ping(1M)`.

Examples—Checking the Response of Hosts on the Network

The simplest version of `ping` sends a single packet to a host on the network. If `ping` receives the correct response, it prints the message *host is alive*.

```
$ ping elvis
elvis is alive
```

With the `-s` option, `ping` sends one datagram per second to a host. It then prints each response and the time that was required for the round trip. For example:

```
$ ping -s pluto
64 bytes from pluto (123.456.78.90): icmp_seq=0. time=10. ms
64 bytes from pluto (123.456.78.90): icmp_seq=5. time=0. ms
64 bytes from pluto (123.456.78.90): icmp_seq=6. time=0. ms
^C
----pluto PING Statistics----
8 packets transmitted, 8 packets received, 0% packet loss

round-trip (ms) min/avg/max = 0/2/10
```

▼ How to Send Packets to Hosts on the Network

Test the reliability of your packet sizes with the `spray` command.

```
$ spray [ -c count -d interval -l packet_size] hostname
```

<code>-i <i>count</i></code>	Number of packets to send.
<code>-d <i>interval</i></code>	Number of microseconds to pause between sending packets. If you don't use a delay, you might run out of buffers.
<code>-l <i>packet_size</i></code>	Is the packet size.
<code><i>hostname</i></code>	Is the system to send packets.

For more information about this command, see `spray(1M)`.

Example—Sending Packets to Hosts on the Network

The following example sends 100 packets to a host (`-c 100`) with each packet having a size of 2048 bytes (`-l 2048`). The packets are sent with a delay time of 20 microseconds between each burst (`-d 20`).

```
$ spray -c 100 -d 20 -l 2048 pluto
sending 100 packets of length 2048 to pluto ...
no packets dropped by pluto
279 packets/sec, 573043 bytes/sec
```

▼ How to Capture Packets From the Network

To capture packets from the network and trace the calls from each client to each server, use `snoop`. This command provides accurate timestamps that allow some network performance problems to be isolated quickly. For more information, see `snoop(1M)`.

```
# snoop
```

Dropped packets could be caused by insufficient buffer space or an overloaded CPU.

▼ How to Check the Network Status

Display network status information, such as statistics about the state of network interfaces, routing tables, and various protocols, with the `netstat` command.

```
$ netstat [-i] [-r] [-s]
```

-i	Displays the state of the TCP/IP interfaces
-r	Displays the IP routing table
-s	Displays statistics for the UDP, TCP, ICMP, and IGMP protocols

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

Examples—Checking the Network Status

The following example shows output from the `netstat -i` command, which displays the state of the interfaces that are used for TCP/IP traffic.

```
$ netstat -i
Name  Mtu  Net/Dest   Address      IpKts  Ierrs OpKts  Oerrs Collis Queue
lo0   8232  software   localhost    1280   0     1280   0     0     0
le0   1500  loopback   venus        1628480 0     347070 16    39354  0
```

This display shows how many packets a machine has transmitted and received on each interface. A machine with active network traffic should show both `IpKts` and `OpKts` continually increasing.

Calculate the network collisions rate by dividing the number of collision counts (`Collis`) by the number of out packets (`OpKts`). In the previous example, the collision rate is 11 percent. A network-wide collision rate greater than 5 to 10 percent can indicate a problem.

Calculate the input packet error rate by dividing the number of input errors by the total number of input packets (`Ierrs/IpKts`). The output packet error rate is the number of output errors divided by the total number of output packets (`Oerrs/OpKts`). If the input error rate is high (over 0.25 percent), the host might be dropping packets.

The following example shows output from the `netstat -s` command, which displays the per-protocol statistics for the UDP, TCP, ICMP, and IGMP protocols.

```
UDP
    udpInDatagrams    =196543    udpInErrors        =      0
    udpOutDatagrams    =187820

TCP
    tcpRtoAlgorithm    =      4    tcpRtoMin           =    200
    tcpRtoMax          =   60000    tcpMaxConn          =     -1
    tcpActiveOpens     =   26952    tcpPassiveOpens     =    420
    tcpAttemptFails    =    1133    tcpEstabResets      =      9
    tcpCurrEstab       =      31    tcpOutSegs          =3957636
    tcpOutDataSegs     =2731494    tcpOutDataBytes     =1865269594
    tcpRetransSegs     =    36186    tcpRetransBytes     =3762520
    tcpOutAck           =1225849    tcpOutAckDelayed    =165044
    tcpOutUrg          =      7    tcpOutWinUpdate     =    315
    tcpOutWinProbe     =      0    tcpOutControl        =   56588
```

Composed September 5, 2002

tcpOutRsts	=	803	tcpOutFastRetrans	=	741
tcpInSegs	=	4587678			
tcpInAckSegs	=	2087448	tcpInAckBytes	=	1865292802
tcpInDupAck	=	109461	tcpInAckUnsent	=	0
tcpInInorderSegs	=	3877639	tcpInInorderBytes	=	-598404107
tcpInUnorderSegs	=	14756	tcpInUnorderBytes	=	17985602
tcpInDupSegs	=	34	tcpInDupBytes	=	32759
tcpInPartDupSegs	=	212	tcpInPartDupBytes	=	134800
tcpInPastWinSegs	=	0	tcpInPastWinBytes	=	0
tcpInWinProbe	=	456	tcpInWinUpdate	=	0
tcpInClosed	=	99	tcpRttNoUpdate	=	6862
tcpRttUpdate	=	435097	tcpTimRetrans	=	15065
tcpTimRetransDrop	=	67	tcpTimKeepalive	=	763
tcpTimKeepaliveProbe	=	1	tcpTimKeepaliveDrop	=	0

IP

ipForwarding	=	2	ipDefaultTTL	=	255
ipInReceives	=	11757234	ipInHdrErrors	=	0
ipInAddrErrors	=	0	ipInCksumErrs	=	0
ipForwDatagrams	=	0	ipForwProhibits	=	0
ipInUnknownProtos	=	0	ipInDiscards	=	0
ipInDelivers	=	4784901	ipOutRequests	=	4195180
ipOutDiscards	=	0	ipOutNoRoutes	=	0
ipReasmTimeout	=	60	ipReasmReqds	=	8723
ipReasmOKs	=	7565	ipReasmFails	=	1158
ipReasmDuplicates	=	7	ipReasmPartDups	=	0
ipFragOKs	=	19938	ipFragFails	=	0
ipFragCreates	=	116953	ipRoutingDiscards	=	0
tcpInErrs	=	0	udpNoPorts	=	6426577
udpInCksumErrs	=	0	udpInOverflows	=	473
rawipInOverflows	=	0			

ICMP

icmpInMsgs	=	490338	icmpInErrors	=	0
icmpInCksumErrs	=	0	icmpInUnknowns	=	0
icmpInDestUnreachs	=	618	icmpInTimeExcds	=	314
icmpInParmProbs	=	0	icmpInSrcQuenchs	=	0
icmpInRedirects	=	313	icmpInBadRedirects	=	5
icmpInEchos	=	477	icmpInEchoReps	=	20
icmpInTimestamps	=	0	icmpInTimestampReps	=	0
icmpInAddrMasks	=	0	icmpInAddrMaskReps	=	0
icmpInFragNeeded	=	0	icmpOutMsgs	=	827
icmpOutDrops	=	103	icmpOutErrors	=	0
icmpOutDestUnreachs	=	94	icmpOutTimeExcds	=	256
icmpOutParmProbs	=	0	icmpOutSrcQuenchs	=	0
icmpOutRedirects	=	0	icmpOutEchos	=	0
icmpOutEchoReps	=	477	icmpOutTimestamps	=	0
icmpOutTimestampReps	=	0	icmpOutAddrMasks	=	0
icmpOutAddrMaskReps	=	0	icmpOutFragNeeded	=	0
icmpInOverflows	=	0			

IGMP:

```

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received

```

```

0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

The following example shows output from the `netstat -r` command, which displays the IP routing table.

```

Routing Table:
  Destination      Gateway      Flags  Ref    Use  Interface
-----
localhost          localhost    UH      0   2817   lo0
earth-bb           pluto       U       3  14293   le0
224.0.0.0          pluto       U       3     0   le0
default            mars-gate   UG      0  14142

```

The fields in the `netstat -r` report are described in Table 46-2.

TABLE 46-2 Output From the `netstat -r` Command

Field Name	Description	
Flags	U	The route is up.
	G	The route is through a gateway.
	H	The route is to a host.
	D	The route was dynamically created by using a redirect.
Ref	Shows the current number of routes that share the same link layer.	
Use	Indicates the number of packets that were sent out.	
Interface	Lists the network interface that is used for the route.	

▼ How to Display NFS Server and Client Statistics

The NFS distributed file service uses a remote procedure call (RPC) facility that translates local commands into requests for the remote host. The remote procedure calls are synchronous. That is, the client application is blocked or suspended until the server has completed the call and returned the results. One of the major factors affecting NFS performance is the retransmission rate.

If the file server cannot respond to a client's request, the client retransmits the request a specified number of times before it quits. Each retransmission imposes system overhead and increases network traffic. Excessive retransmissions can cause network performance problems. If the retransmission rate is high, you could look for the following:

- Overloaded servers that take too much time to complete requests
- An Ethernet interface that is dropping packets
- Network congestion, which slows the packet transmission

Table 46-3 describes the `nfsstat` options to display client and server statistics.

TABLE 46-3 Commands for Displaying Client/Server Statistics

Command	Display
<code>nfsstat -c</code>	Client statistics
<code>nfsstat -s</code>	Server statistics
<code>netstat -m</code>	Network statistics for each file system

Use `nfsstat -c` to show client statistics, and `nfsstat -s` to show server statistics. Use `netstat -m` to display network statistics for each file system. For more information, see `nfsstat(1M)`.

Examples—Displaying NFS Server and Client Statistics

The following example displays RPC and NFS data for the client `pluto`.

```
$ nfsstat -c

Client rpc:
Connection oriented:
calls    badcalls  badxids  timeouts  newcreds  badverfs  timers
1595799  1511      59       297       0         0         0
cantconn nomem     interrupts
1198     0        7
Connectionless:
calls    badcalls  retrans  badxids  timeouts  newcreds  badverfs
80785   3135     25029   193      9543      0         0
timers  nomem     cantsend
17399   0        0

Client nfs:
calls    badcalls  clgets  cltoomany
1640097  3112     1640097  0
Version 2: (46366 calls)
null     getattr  setattr  root      lookup    readlink  read
0 0%     6589 14%  2202 4%  0 0%     11506 24%  0 0%     7654 16%
wrcache  write    create   remove    rename    link      symlink
0 0%     13297 28%  1081 2%  0 0%     0 0%     0 0%     0 0%
mkdir    rmdir    readdir  statfs
24 0%    0 0%    906 1%  3107 6%
Version 3: (1585571 calls)
null     getattr  setattr  lookup    access    readlink  read
0 0%     508406 32%  10209 0%  263441 16%  400845 25%  3065 0%  117959 7%
write    create   mkdir    symlink    mknod     remove    rmdir
```

```

69201 4% 7615 0%    42 0%    16 0%    0 0%    7875 0%  51 0%
rename  link      readdir readdir+  fsstat  fsinfo  pathconf
929 0%  597 0%    3986 0% 185145 11% 942 0%  300 0%  583 0%
commit
4364 0%

```

Client nfs_acl:

Version 2: (3105 calls)

```

null      getacl      setacl      getattr     access
0 0%      0 0%      0 0%      3105 100%  0 0%

```

Version 3: (5055 calls)

```

null      getacl      setacl
0 0%      5055 100%  0 0%

```

The output of the `nfsstat -c` command is described in Table 46–4.

TABLE 46–4 Output From the `nfsstat -c` Command

Field	Description
calls	The total number of calls that were sent.
badcalls	The total number of calls that were rejected by RPC.
retrans	The total number of retransmissions. For this client, the number of retransmissions is less than 1 percent (10 timeouts out of 6888 calls). These retransmissions might be caused by temporary failures. Higher rates might indicate a problem.
badxid	The number of times that a duplicate acknowledgment was received for a single NFS request.
timeout	The number of calls that timed out.
wait	The number of times a call had to wait because no client handle was available.
newcred	The number of times the authentication information had to be refreshed.
timers	The number of times the time-out value was greater than or equal to the specified time-out value for a call.
readlink	The number of times a read was made to a symbolic link. If this number is high (over 10 percent), it could mean that there are too many symbolic links.

The following example shows output from the `nfsstat -m` command.

```

pluto$ nfsstat -m
/usr/man from pluto:/export/svr4/man
Flags: vers=2,proto=udp,auth=unix,hard,intr,dynamic,
       rsize=8192, wsize=8192,retrans=5
Lookups: srvt=13 (32ms), dev=10 (50ms), cur=6 (120ms)
All:      srvt=13 (32ms), dev=10 (50ms), cur=6 (120ms)

```


This output of the `nfsstat -m` command, which is displayed in milliseconds, is described in Table 46-5.

TABLE 46-5 Output From the `nfsstat -m` Command

Field	Description
srtt	The smoothed average of the round-trip times
dev	The average deviations
cur	The current “expected” response time

If you suspect that the hardware components of your network are creating problems, you need to look closely at the cabling and connectors.

Composed September 5, 2002

APPENDIX **A**

System Administration Guide: Resource Management and Network Services Updates

Solaris 9 9/02 Updates

Documentation covering the use of the IPQoS flow accounting module in conjunction with extended accounting to capture network flow statistics was added to Chapter 7.

Solaris 9 12/02 Updates

Documentation covering multiple instance support in NCA was added to Chapter 2.

Composed September 5, 2002

Glossary

asppp	A version of PPP that was included with the Solaris operating environment from the Solaris 2.4 to the Solaris 8 releases. asppp supported asynchronous PPP communications only.
authentication	The act of verifying the identity that is supplied over the network by a remote user or entity, such as a program. Some authentication protocols enable you to build databases of authentication credentials from potential users. Other authentication protocols use certificate chains of trust that are generated by a certificate authority for authentication purposes. These credentials can authenticate users when they try to communicate with you or use your site's services.
asynchronous PPP	A form of PPP that runs over asynchronous serial lines, which transfer data one character at a time. The most common form of PPP configuration, the dial-up link, uses asynchronous PPP communications.
broadcast	A data-link layer procedure that is used to transmit packets to every machine on a subnet. Broadcast packets are typically not routed beyond the subnet.
Callback Control Protocol (CBCP)	A proprietary Microsoft PPP extension that is used to negotiate a callback session. Solaris PPP 4.0 supports only the client (initial caller) side of this protocol.
Compression Control Protocol (CCP)	A subprotocol of PPP that negotiates the use of data compression on the link. Unlike header compression, CCP compresses all data within packets that are sent on the link.
challenge-handshake authentication protocol (CHAP)	An authentication protocol that can be used to verify the identity of a caller on a PPP link. CHAP authentication uses the notion of the <i>challenge</i> and <i>response</i>, where the machine that receives a call challenges the caller to prove its identity. See also password authentication protocol (PAP).

channel service unit (CSU)	A synchronous telecommunications device that provides a local interface to a leased telecommunications line and terminates that line. In the United States, a CSU terminates a T1 line and provides a DS1 or DSX interface. Internationally, the CSU is typically owned by the telephone company provider. See also CSU/DSU and data service unit (DSU).
CHAP secret	An ASCII or binary string that is used for identification purposes and is known to both peers on a PPP link. The CHAP secret is stored in clear text in a system's <code>/etc/ppp/chap-secrets</code> file but is never sent over the PPP link, not even in encrypted form. The CHAP protocol verifies that a hash of the CHAP secret that is used by a caller matches a hash of the CHAP secret entry for the caller in the recipient's <code>/etc/ppp/chap-secrets</code> file.
chat script	Instructions that tell a modem how to establish a communications link between itself and a remote peer. Both the PPP and UUCP protocols use chat scripts for establishing dial-up links and dial-back calling.
CSU/DSU	A synchronous telecommunications device that combines the CSU and DSU devices and is used on a leased-line PPP link. The CSU/DSU translates signals from a peer to a leased line. Most CSU/DSUs do not need a chat script to establish the link. CSU/DSUs are often configured by the leased-line provider. See also channel service unit (CSU) and data service unit (DSU).
data service unit (DSU)	A synchronous telecommunications device that is used on a leased-line PPP link. The DSU converts between data-framing formats that are used on telecommunications lines and provides a standard data communications interface. See also channel service unit (CSU) and CSU/DSU.
dial-in server	The peer that negotiates and establishes the recipient end of a dial-up PPP link after receiving a call from a dial-out machine. Though the term "dial-in server" is in common use, the dial-in server does not function in accordance with the client-server paradigm. Rather, it is simply the peer that responds to the request to set up a dial-up link. After it is configured, a dial-in server can receive calls from any number of dial-out machines.
dial-out machine	The peer that initiates the call to establish a dial-up PPP link. After it is configured, the dial-out machine can call any number of dial-in servers. The dial-out machine typically provides authentication credentials before the dial-up link can be established.
dial-up PPP link	A PPP connection that involves a peer and a modem at either end of a telephone line or similar communications medium, such as a medium that is provided by ISDN. The term "dial-up" refers to the sequence in

	link negotiation when the local modem dials up the remote peer by using the peer's telephone number. The dial-up link is the most common and least expensive PPP configuration.
Directory Agent (DA)	Optional SLP agent that stores and maintains a cache of service advertisements that are sent by the service agent (SA). When deployed, the DA resolves user agent (UA) service requests. The DA responds to active solicitations from the SA and UA for directory advertisements. As a result, the SA and UA discover the associated DAs and <i>scopes</i> . A DA sends periodic unsolicited advertisements through which UAs and SAs discover the DA within shared scopes.
expect-send	A scripting format that is used in PPP and UUCP chat scripts. The chat script begins with the text or instruction to <i>expect</i> from the remote peer. The next line contains the response to be <i>sent</i> from the local host after it receives the correct expect string from the peer. Subsequent lines repeat the expect-send instructions between local host and peer until all instructions that are required to establish communications are successfully negotiated.
extended accounting	A flexible way to record resource consumption on a task or process basis in the Solaris operating environment.
fair share scheduler	A scheduling class that allows you to allocate CPU time that is based on shares. Shares define the portion of the system's CPU resources that are allocated to a project.
Internet Protocol Control Protocol (IPCP)	A subprotocol of PPP that negotiates the IP addresses of the peers on the link. IPCP also negotiates header compression for the link and enables the use of the network layer protocols.
Internet Protocol Version 6 Control Protocol (IPV6CP)	See Internet Protocol Control Protocol (IPCP).
ISDN terminal adaptor (TA)	A signal-adapting device that provides a modem-like interface for a dial-up PPP link over an ISDN network. You use the same Solaris PPP 4.0 configuration files to configure an ISDN TA as you use for configuring a standard modem.
leased-line PPP link	A PPP connection that involves a host and a CSU/DSU that are connected to a synchronous network medium leased from a provider. OC3 and T1 are common examples of leased-line media. Though easier to administer, leased-line links are more expensive than dial-up PPP links and, therefore, are less common.
legacy services	A networked service that is not SLP-enabled. You can create a proxy registration to register a legacy service with SLP. SLP-based clients can then discover legacy services (see Chapter 21).

link	In PPP, the communications connection that is negotiated and established between two peers. Solaris PPP 4.0 supports two types of links: dial up and leased line.
link control protocol (LCP)	A subprotocol of PPP that is used to negotiate the initial set of link parameters between the peers. Part of the function of LCP is to test the link integrity, so many link-related problems manifest themselves as LCP failure.
Microsoft CHAP (MS-CHAP)	A proprietary Microsoft authentication protocol for PPP. Solaris PPP 4.0 supports versions 1 and 2 of this protocol in both client and server mode.
multicast	A network layer procedure that is used to send datagram packets to multiple machines on an IP network. Packets are not handled by every machine as is the situation with broadcast routing. Multicast requires that routers be configured with special routing protocols.
password authentication protocol (PAP)	An authentication protocol that can be used to verify the identity of a caller on a PPP link. PAP uses a cleartext password that is passed over the link, which makes it possible to store the password on one of the endpoint machines. For example, PAP can use the login and password entries in the UNIX <code>passwd</code> database on the machine that receives a call to verify the identity of the caller. See also challenge-handshake authentication protocol (CHAP).
peer	In PPP, an individual computer at one end of a PPP communications link, which consists of two peers that are connected by communications media. You can configure many types of computing equipment as a peer, such as a workstation, personal computer, router, or mainframe.
point-to-point protocol (PPP)	A data-link layer protocol that provides a standard method for transferring datagrams over point-to-point media. A PPP configuration consists of two endpoint computers called <i>peers</i>, and the telephone lines or another bidirectional link that the peers use for communication. The hardware and software connection between the two peers is considered the <i>PPP link</i>. PPP is composed of a number of subprotocols, including PAP, CHAP, LCP, and CCP. Numerous PPP implementations are available. The Solaris 9 operating environment includes Solaris PPP 4.0.
pool	In resource management, a configuration mechanism that is used to partition machine resources.
PPP over Ethernet (PPPoE)	A proprietary protocol from RedBack Networks that enables hosts to run PPP sessions over an Ethernet link. PPPoE is commonly used with Digital Subscriber Line (DSL) services.
project	A network-wide administrative identifier for related work.

resource	In resource management, an aspect of the computing system that can be manipulated with the intent to change application behavior.
resource control	In resource management, a per-process resource limit that is extended to the task and project entities.
resource management	A functionality that enables you to control how applications use available system resources.
scope	A grouping of UAs and SAs that are arranged administratively, topologically, or in some other manner. You can use scopes to modify how you provision access to services across the enterprise.
service advertisements	Information that is distributed by an SA that describes a service. A service advertisement consists of a URL and a collection of attribute/value list pairs that describe a service. All service advertisements have a lifetime. After the lifetime expires, a service advertisement is no longer valid unless reregistered.
Service Agent (SA)	The SLP agent that maintains service advertisements for networked services. If no DA is available, the SA answers multicast service requests from UAs. If a DA is available, the SA registers and, optionally, deregisters services with DAs that support its scopes.
service URL	A URL that is used to advertise the network location of services. The URL contains the service type, host name, or network address of the service host. The URL might also contain a port number and other information that is required to use the service.
SLP daemon (slpd)	The daemon process that acts as a DA or an SA server in the Solaris implementation of SLP. Service processes on the host register service advertisements with <code>slpd</code> instead of maintaining the advertisements individually. Each process contains an SA client library that communicates with <code>slpd</code> when the daemon is configured as the SA server. The SLP daemon forwards all registrations and deregistrations to DAs. The daemon times out expired service advertisements and maintains a table of the available DAs by performing active and passive DA discovery. Through such mechanisms, DA information is provided to UA clients. UA clients use <code>slpd</code> on a host only for DA information. The SLP daemon is installed on a host as part of the Solaris 9 operating environment. You can optionally configure <code>slpd</code> as a DA.
synchronous PPP	A form of PPP that runs over synchronous digital lines, which transfer data as a continuous stream of raw bits. The leased-line PPP link uses synchronous PPP.
task	In resource management, a process collective that represents a set of work over time. Each task is associated with one project.

Composed September 5, 2002

trusted callers	In PPP, remote peers that a dial-in server grants access to by including the peers' security credentials in the server's PAP or CHAP secrets database.
User Agent (UA)	The SLP agent that acts on behalf of the user application. The agent queries for the identity of corresponding scopes, directory agents, and service advertisements.

Index

Numbers and Symbols

- /
- /- as master map mount point, 233, 236
- #
 - comments in direct maps, 235
 - comments in indirect maps, 236
 - comments in master map (auto_master), 233
- /
- master map names preceded by, 233
- root directory
 - mounting by diskless clients, 145
- & (ampersand), in autofs maps, 249
- * (asterisk), in autofs maps, 250
- (dash)
 - dial-code abbreviation, 590
 - in autofs map names, 245
 - Line2 field placeholder, 595
 - Speed field placeholder, 590
- .(dot)
 - rcp command syntax, 676, 677
- + (plus sign)
 - /etc/hosts.equiv file syntax, 659, 660
 - in autofs map names, 245, 246
- ~ (tilde)
 - abbreviated path names, 673, 674
 - rcp command syntax, 676, 677
- = dial-code abbreviation, 590
- \ in maps, 233, 235, 236
- (DAs) SLP
 - dial-up networks discovery, 271
 - heartbeat, 271
- \$OPENWINHOME/bin/mailtool
 - command, 370

A

- a option
 - showmount command, 216
 - umount command, 208
- Ac option, sendmail command, 392
- access control list (ACL), NFS and, 147
- access server (PPP)
 - commands and files for configuring, 557, 558
 - configuring, for PPPoE, 501, 502, 504, 560
 - definition, 437
 - /etc/ppp/chap-secrets file, 562
 - /etc/ppp/options file, 561
 - /etc/ppp/pap-secrets file, 561
 - planning task map, 455
 - restricting an interface to PPPoE clients, 503
 - task map for configuring, 498
- acctadm command, 85, 87
- activating extended accounting, 85
- ACU keyword of Type field, 594
- address assignment
 - PPP, 552, 553, 554
- administering resource pools, 117
- administrative commands (UUCP), 573
- administrative files (UUCP), 618, 619
 - cleanup, 580
 - execute files (X.), 573, 619
 - lock files (LCK), 618
 - temporary data files (TM), 618
 - work files (C.), 618, 619
- AdminTool's Database Manager
 - administering mail aliases, 327, 377
 - administering NIS+ mail aliases, 379

- aliasadm command, 366
- aliases
 - creating, 361, 362
 - definition, 361
 - /etc/mail/aliases file, 377, 378
 - loops, 346
 - NIS+ mail_aliases table, 379
 - NIS aliases map, 378
 - verifying, 345
- aliases.db file, 334, 366
- aliases.dir file, 334, 366
- aliases file, 366, 582
- aliases.pag file, 334, 367
- ALL value in COMMANDS option, 610
- already mounted message, 190
- alternative commands, sendmail
 - command, 353
- Am option, sendmail command, 392
- ampersand (&), in autofs maps, 249
- anon option, share command, 211
- anonymous ftp
 - accounts, 666
 - setting up, 640
- Any keyword
 - Grades file (UUCP), 615, 616
 - Speed field (UUCP), 589
- Any Time field entry, 588
- applications, hung, 194
- ARCH map variable, 245
- asppp
 - See* asynchronous PPP (asppp)
- asppp links
 - configuring
 - UUCP databases, 576
- asppp2pppd conversion script
 - converting to Solaris PPP 4.0, 568
 - standard asppp configuration, 565
 - viewing files converted to Solaris PPP 4.0, 569
- ASSERT ERROR, 622
- ASSERT error messages (UUCP), 585, 619, 621
- asterisk (*), in autofs maps, 250
- asynchronous PPP (asppp)
 - configuring UUCP databases, 575
 - converting to Solaris PPP 4.0, 568
 - difference with Solaris PPP 4.0, 426
 - documentation, 426
 - files in a configuration, 565
 - asyncmap option (PPP), 530
 - audio files, mailbox space requirements
 - and, 364
 - Australian National University (ANU) PPP,
 - compatibility with Solaris PPP 4.0, 426
 - auth option (PPP), 486
 - authenticatee (PPP), 435
 - authentication
 - See also* authentication (PPP)
 - DH, 230, 231
 - fixing common problems, 519
 - remote logins using ftp command, 666, 667, 668
 - remote logins using rlogin command, 658, 661, 665
 - direct or indirect logins, 661
 - /etc/hosts.equiv file, 659, 660
 - network or remote system
 - authentication, 658, 660, 661
 - .rhosts files, 660, 661
 - RPC, 230
 - UNIX, 229, 230
 - authentication (PPP)
 - authenticatee, 435
 - authenticator, 435
 - configuring CHAP, 491, 492, 493, 494
 - See also* Challenge-Handshake Authentication Protocol (CHAP)
 - configuring PAP
 - See also* Password Authentication Protocol (PAP)
 - default policy, 435
 - example of CHAP, 452
 - example of PAP, 450
 - planning, 449, 450, 451, 452, 453
 - prerequisites before configuring, 449
 - process diagram
 - for CHAP, 550
 - for PAP, 547
 - secrets file, 435
 - for PAP, 484
 - support for leased lines, 435
 - task maps for configuring, 481, 482, 490
 - trusted callers, 435
 - authenticator (PPP), 435
 - auto_direct file, 317
 - auto_home map
 - /home directory, 174

- auto_home map (Continued)
 - /home directory server setup, 175
 - /home mount point, 232, 233
 - autofs, 250
 - administering maps, 170
 - browsability, 149, 180
 - consolidating project-related files, 176
 - default behavior, 247, 249
 - features, 149
 - /home directory, 174
 - home directory server setup, 175
 - maps
 - administrative tasks, 249
 - browsability and, 149
 - cachefs option, 174
 - CD-ROM file system, 173
 - default behavior, 247, 249
 - direct, 234, 235
 - hsfs option, 173
 - indirect, 236, 237
 - master, 232, 233
 - modifying, 247
 - network navigation, 240
 - PC-DOS file system, 173
 - pcfs option, 173
 - read-only file selection, 242, 244
 - referring to other maps, 245, 246
 - starting the navigation process, 233, 240
 - types, 170
 - variables, 244, 245
 - metacharacters, 249, 250
 - mount process, 240, 241
 - mounting file systems, 159
 - name service use, 248
 - name space data, 149
 - NFS URL and, 180
 - non-NFS file system access, 172, 173
 - operating systems
 - supporting incompatible versions, 178
 - overview, 145
 - public file handle and, 180
 - reference, 249, 250
 - replicating shared files across several servers, 179
 - shared name space access, 177
 - special characters, 250
 - starting, 163
 - stopping, 163
 - autofs (Continued)
 - task map, 168
 - troubleshooting, 189
 - unmount process, 242
 - autofs script, 239
 - starting autofs, 163
 - stopping autofs, 163
 - Automatic Call Unit (ACU)
 - Devices file Type field, 594
 - troubleshooting, 583
 - UUCP hardware configuration, 571
 - automatic file-system sharing, 152, 153
 - automatic mounting
 - /var/mail directory, 317, 363
 - automount command
 - autofs and, 145
 - error messages, 189
 - modifying autofs master map (auto_master), 171
 - modifying direct autofs map, 171
 - overview, 238
 - v option, 189
 - when to run, 170
 - automountd daemon
 - autofs and, 145
 - description, 149
 - mounting and, 149
 - nobrowse option, 181
 - overview, 238
 - AutoRebuildAliases option, sendmail command, 400
- B**
- b escape character
 - Dialers file, 601
 - Systems file chat script, 591
 - background file mounting option, 206
 - backslash (\) in maps, 233, 235, 236
 - backslash escape character
 - Dialers file send strings, 600
 - Systems file chat script, 591
 - backspace escape character, 591, 601
 - backups, mail servers and, 364
 - bad argument specified with index option, 192
 - bad key message, 189

BAD LINE message, 620
BAD LOGIN/MACHINE COMBINATION message, 621
BAD LOGIN_UID message, 620
BAD OPTION message, 620
bad pooladm configuration, recovery from, 123
BAD SPEED message, 620
BAD UID message, 620
BadRcptThrottle option, sendmail command, 393
bg option, mount command, 206
binding to a resource pool, 124
booting
 diskless client security, 231
 mounting file systems, 157
-bP option, sendmail command, 392
Break escape character, 592
 Dialers file, 601
 Systems file chat script, 591, 592
broadcast (SLP), 276, 284, 288
browsability
 disabling, 180
 overview, 149
browsing, with an NFS URL, 167
buffer parameter,
 /etc/default/nfslog.conf file, 199
bye command (FTP), 668

C

C. UUCP work files
 cleanup, 580
 description, 618, 619
c escape character
 Dialers file, 601
 Systems file chat script, 591
-c option, nfsd daemon, 202
cache and NFS version 3, 146
cache file system type
 autofs access using, 173, 174
cachefs option, autofs maps, 174
call option (PPP), calling a dial-in server, 473
callback
 enabling dialback through chat script, 592
 Permissions file option, 609

CALLBACK option of Permissions file, 608, 609
CALLBACK REQUIRED message, 622
CALLER SCRIPT FAILED message, 622
CAN'T ACCESS DEVICE message, 622
CAN'T ALLOCATE message, 620
CAN'T CHDIR message, 620
CAN'T CHMOD message, 620
CAN'T CLOSE message, 620
CAN'T CREATE message, 620
CAN'T FORK message, 621
CAN'T LINK message, 620
CAN'T LOCK message, 620
can't mount message, 189
CAN'T MOVE TO CORRUPTDIR message, 620
CAN'T OPEN message, 620
CAN'T READ message, 620
CAN'T STAT message, 620
CAN'T UNLINK message, 620
CAN'T WRITE message, 620
canceling, remote logins, 658
cannot receive reply message, 191
cannot send packet message, 191
cannot use index option without public option message, 192
carriage-return escape characters, 591, 592, 601
CD-ROM applications, accessing with
 autofs, 173
cfsadmin command, accessing NFS file systems, 174
Challenge-Handshake Authentication Protocol (CHAP)
 authentication process, 550, 552
 definition, 549
 example configuration, 452
 syntax of /etc/ppp/chap-secrets, 550
 task maps for configuring, 490
changing
 /etc/shells file, 343
 .forward file search path, 342
changing resource controls temporarily, 95
CHAP credentials database
 creating
 for a dial-in server, 491
 for trusted callers, 493

- chat program in PPP
 - See also* chat script
- chat script
 - creating an executable chat program, 545
 - designing the chat script, 536
 - examples (PPP)
 - basic modem chat script, 537
 - for an ISDN TA, 542, 543
 - script for calling an ISP, 539
 - UNIX-style login chat script, 540, 541
 - examples for PPP
 - UNIX-style login chat script, 464
 - invoking, in PPP, 544
 - UUCP chat script field, 590
- Chat Script field, /etc/uucp/Systems file, 590
- check_eoh rule set, sendmail command, 416
- check_etrn rule set, sendmail command, 417
- check_expn rule set, sendmail command, 417
- check-hostname script, 319, 321, 369
- check-permissions script, 369
- check_vrfy rule set, sendmail command, 417
- chkey command, enabling secure NFS, 164
- Class field, Devices file, 595
- clear_locks command, 205
- client-side failover
 - enabling, 160
 - NFS locking and, 226
 - NFS support, 147
 - overview, 224
 - replicated file systems, 225
 - terminology, 225
- clientmqueue directory, 370
- ClientPortOption option, sendmail command, 393
- ClientPortOptions option, sendmail command, 400
- clients
 - See* mail clients, NFS clients, NTP client, and PPPoE client
 - displaying information about, 681, 687, 689
 - tracing calls to servers, 681, 683
- CLOCAL flag, turning on and off, 591
- closing remote system connections, 668
- collision rate (network), 684
- command-line options, sendmail command, 391
- commands
 - execute (X.) UUCP files, 573, 619
 - extended accounting, 84
 - hung programs, 194
 - NFS commands, 204
 - remote execution using UUCP, 606, 609, 612
 - UUCP troubleshooting, 585
- COMMANDS option of Permissions file, 609, 610, 613
- VALIDATE option, 610, 612
- comments
 - in direct maps, 235
 - in indirect maps, 236
 - in master map (auto_master), 233
- compat_check FEATURE() declaration, 409
- compilation flags, sendmail command, 352
- conf_FORWARD_PATH definition, 342, 343
- configuration examples for PPP
 - CHAP authentication, 452
 - dial-up link, 443
 - leased-line link, 447
 - PAP authentication, 450
 - PPPoE tunnel, 456
- configuration file options, sendmail command, 392
- configuration files
 - sendmail command, 376
 - UUCP, 613
- configuration tasks for PPP
 - authentication, 481
 - diagnosing configuration problems, 513
 - dial-up link, 459
 - leased lines, 475
 - PPPoE tunnel, 497
- configuring
 - asppp links
 - UUCP databases, 575, 576
 - mail gateways, 364
 - UUCP
 - adding logins, 578
 - database files, 575, 576
 - shell scripts, 579, 581
 - TCP/IP networks, 581
- configuring resource controls, 90
- connect option (PPP)
 - /etc/ppp/peers/peer-name file, 533

- connect option (PPP) (Continued)
 - example, 466
 - running the pppoecc utility, 563
 - to invoke a chat script, 543
- ConnectionRateThrottle option,
 - sendmail command, 393
- consolidating project-related files, 176
- ControlSocketName option, sendmail command, 393
- CONVERSATION FAILED message, 621
- conversation key, 231
- copying files (remote)
 - using ftp, 667
 - using rcp, 673, 677
- could not start daemon message, 192
- could not use public filehandle message, 192
- couldn't create mount point message, 189
- CPU map variable, 245
- CPU share configuration, 106
- creating
 - /etc/shells file, 343
 - keyed map file, 335
 - postmaster alias, 335
 - postmaster mailbox, 336
- creating resource pools, 119
- credentials
 - CHAP authentication, 491
 - description, 230
 - PAP authentication, 483
 - UNIX authentication, 230
- crontab file, for UUCP, 579
- crtsets option (PPP), 464
- CSU/DSU
 - configuring, 477
 - definition, 433
 - fixing common problems, 518
- cu command
 - checking modems or ACUs, 583
 - description, 574
 - multiple or different configuration files, 575, 604
 - printing Systems lists, 605
- current user, 674
- CYCLE_FREQUENCY parameter,
 - /etc/default/nfslogd file, 198

D

- D. UUCP data files, cleanup, 580
- D escape character
 - Devices file, 598
 - Dialers file, 598
- d escape character
 - Dialers file, 601
 - Systems file chat script, 591
- d option
 - cu command, 583
 - showmount command, 216
- DA (SLP)
 - advertising, 268
 - discovery, 268
- DA_BUSY_NOW, 286
- DA discovery (SLP), 278
- DA heartbeat, frequency, 268
- daemon running already message, 192
- DaemonPortOptions option
 - sendmail command, 393, 401
- daemons
 - automountd
 - autofs and, 145
 - overview, 238
 - lockd, 201
 - mountd
 - checking response on server, 185
 - enabling without rebooting, 187
 - not registered with rpcbind, 193
 - remote mounting requirement, 182
 - verifying if running, 186, 193
- nfsd
 - checking response on server, 184
 - description, 202
 - enabling without rebooting, 187
 - remote mounting requirement, 182
 - syntax, 202
 - verifying if running, 186
- nfslogd, 203
- required for remote mounting, 182
- rpcbind
 - dead or hung, 193
 - mountd daemon not registered, 193
- statd, 203
- DAs (SLP), 284, 286, 288
 - advertising, 268, 269, 271
 - deploying, 272, 284
 - dial-up networks discovery, 269, 270

- DAs (SLP) (Continued)
 - discovery, 268, 269, 272, 282
 - eliminating multicast, 269
 - heartbeat, 271, 273
 - multicast, 272
 - removing, 271
- dash (-)
 - dial-code abbreviation, 590
 - in autofs map names, 245
 - Line2 field placeholder, 595
 - Speed field placeholder, 590
- data (D.) UUCP files, cleanup, 580
- DataFileBufferSize option, sendmail
 - command, 394
- date, synchronizing with another system, 57
- day entries for Time field, 588
- ddd escape character, Systems file chat script, 592
- deactivating extended accounting, 87
- DeadLetterDrop option, sendmail
 - command, 394
- debug option for PPP, 508
- debugging
 - UUCP transmissions, 583, 584
- debugging PPP
 - debugging chat scripts, 515
 - diagnosing configuration problems, 513
 - diagnosing network problems, 509
 - diagnosing PPPoE problems, 519
 - diagnosing serial line problems, 517
 - fixing communications problems, 511
 - fixing modem problems, 514
 - turning on debugging, 508
- default file-system type, 197
- default keyword of User-job-grade field, 615
- default project, 70
- defaultdir parameter,
 - /etc/default/nfslog.conf file, 199
- defaulttroute option (PPP), 533, 564
- delay_checks FEATURE() declaration, 409
- delay escape character, 591, 601
- DelayLA option, sendmail command, 394
- deleting, .rhosts files, 661
- DeliverByMin option, sendmail
 - command, 394
- delivery agent flags, sendmail command, 412
- demand initialization script for PPP, 479
- desktop-publishing files, mailbox space requirements and, 364
- /dev/nca file, NCA and, 50
- Devconfig file
 - description, 574, 616
 - format, 616
- DEVICE FAILED message, 622
- DEVICE LOCKED message, 621
- device transmission protocols, 598, 599
- device type for UUCP communication link, 589
- Devices file
 - Class field, 595
 - description, 574, 575, 593
 - Dialer-Token-Pairs field, 596, 598
 - format, 593
 - Line2 field, 595
 - Line field, 595
 - multiple or different files, 604
 - protocol definitions, 598, 599
 - Systems file Speed field and, 590
 - Systems file Type field and, 594
 - Type field, 594, 595
- dfstab file
 - automatic file-system sharing, 153
 - disabling mount access for one client, 160
 - enabling NFS server logging, 155
 - enabling secure NFS, 165
 - enabling WebNFS service, 154
 - secure NFS option, 165
 - syntax for NFS file systems, 153
- DH authentication
 - dfstab file option, 165
 - overview, 230, 231
 - password protection, 229
 - secure NFS and, 163
 - user authentication, 229
- diagnostics for PPP, 507, 508
 - leased-line link, 508
 - log file for a PPPoE tunnel, 519
 - turning on, with pppd, 507
- dial-code abbreviations, 575, 590
- DIAL FAILED message, 621
- dial-in server
 - configuring
 - CHAP authentication, 491, 492
 - modem, 468
 - PAP authentication, 483, 484, 485, 486
 - serial line communications, 471, 529

- configuring (Continued)
 - serial port, 468
 - creating accounts for PPP users, 470, 531
 - definition, 429
 - planning information, 443, 469
 - receiving calls, 473
 - task map for configuring, 467
 - UUCP, 592
- dial-out machine
 - addressing
 - dynamic, 552
 - static, 553
 - calling the remote peer, 473
 - configuring
 - CHAP authentication, 493, 494
 - connection with a peer, 465
 - modem, 462
 - PAP authentication, 487
 - serial line communications, 463
 - serial port, 462
 - configuring a serial line with
 - `/etc/ppp/options.ttyname`, 530
 - creating a chat script, 464
 - definition, 429
 - planning information, 442
 - task map for configuring, 460
 - using `.ppprc` in the `$HOME` directory, 532
- dial-up link, 507
 - authentication for the link, 435
 - creating chat scripts, 536, 537, 538, 539, 540, 541, 542, 543
 - definition, 429
 - diagnosing common problems, 509, 510, 511, 512, 513, 514, 517
 - dial-up process, 431
 - example, 443
 - initiating a call to a peer, 473
 - modem settings, basic, 462
 - parts of the link, 430
 - planning, 442, 443, 444
 - task map, 459
 - templates for configuration files, 461
- dialback
 - CALLBACK option of Permissions file, 608, 609
 - enabling through chat script, 592
- Dialcodes file, 575
 - description, 603
- Dialer-Token-Pairs field of Devices file, 596, 597, 598
- Dialers file, 599, 602
 - code example, 600
 - description, 575, 599
 - Devices file DTP field and, 597
- Digital Subscriber Line Access Multiplexer (DSLAM), for PPPoE, 439
- `dir` must start with `''` message, 190
- direct I/O mounting option, 206
- direct keyword of DTP field, 596
- Direct keyword of Type field, 594
- direct link UUCP configuration, 571
- direct maps (autofs)
 - comments in, 235
 - description, 170
 - example, 234
 - modifying, 171
 - overview, 235
 - syntax, 235
 - when to run automount command, 170
- direct remote logins
 - indirect logins or
 - `rlogin` command, 661
 - using `rlogin` command, 665
- directories (UUCP)
 - administration, 573
 - error messages, 585
 - public directory maintenance, 583
- directory agent (SLP), 253, 268, 272, 285, 286
- DirectSubmissionModifiers option,
 - `sendmail` command, 394
- disabling
 - autofs browsability
 - overview, 180
 - tasks, 180
 - `.forward` files, 342
 - large file creation, 159
 - mount access for one client, 160
 - NCA, 48
 - NCA logging, 49
- discovery requests (SLP), 277
- diskless clients
 - manual mounting requirements, 145
 - security during boot process, 231
- displaying extended accounting status, 85
- displaying network information, 681, 682, 683, 689

- DNS name service, sendmail program
 - and, 322
- dnsbl FEATURE() declaration, 410, 412
- domain directory, 368
- domain names, Secure NFS system and, 163
- domains
 - definition, 163
 - remote logins and, 658
 - subdomains and, 358
- DontBlameSendmail option, sendmail
 - command, 395
- DOS files, accessing with autofs, 173
- dot (.)
 - in domain addresses, 359
 - in mailbox names, 360
 - rcp command syntax, 676, 677
- DoubleBounceAddress option, sendmail
 - command, 395
- drift file, 58
- dropped packets, 683
- DSL
 - See PPPoE
- DSL modem, 439
- dtmail mail user agent, 370
- dynamic addressing, PPP, 552

E

- E escape character
 - Dialers file, 601
 - Systems file chat script, 591
- e escape character
 - Dialers file, 601
 - Systems file chat script, 591
- e option, showmount command, 216
- e protocol in Devices file, 598
- echo checking, 591, 601
- editmap command, 370, 420
- email, UUCP maintenance, 582
- enabling
 - automatic file-system sharing, 153
 - client-side failover, 160
 - NCA, 46
 - NCA logging, 49
 - NFS server logging, 155
 - secure NFS system, 164
 - WebNFS service, 154

- enhdnsbl FEATURE() declaration, 410, 412
- entry format, project file, 71
- EOT escape character, 592
- equal sign (=) in dial-code abbreviation, 590
- equates for delivery agents, sendmail
 - command, 413
- error checking message, 193
- error locking message, 192
- error messages
 - generated by automount -v, 189
 - miscellaneous automount messages, 190
 - No such file or directory, 193
 - open errors
 - NFS and, 146
 - Permission denied, 193
 - sendmail program, 348
 - server not responding
 - during mounting, 207
 - hung programs, 194
 - keyboard interrupt for, 182
 - remote mounting problems, 193, 194
 - write errors
 - NFS and, 146
- errors directory (UUCP), 585
- escape characters
 - Dialers file send strings, 600
 - Systems file chat script, 591
- /etc/.rootkey file
 - enabling secure NFS, 164, 165
- /etc/asppp.cf configuration file, 566
- /etc/auto_direct file, 317
- /etc/default/fs file, 197
- /etc/default/nfs file
 - description, 197
 - lockd and, 201
 - nfsd and, 202
- /etc/default/nfslogd file,
 - description, 197
- /etc/default/sendmail file, 370, 381
- /etc/dfs/dfstab file
 - automatic file-system sharing, 153
 - disabling mount access for one client, 160
 - enabling NFS server logging, 155
 - enabling secure NFS, 165
 - enabling WebNFS service, 154
 - secureNFS option, 165
- /etc/dfs/fstypes file, description, 198
- /etc/dfs/sharetab file, description, 198

/etc/hostname.interface file, NCA and, 50
/etc/hosts.equiv file, 659, 660
/etc/hosts file, 51, 312, 313
/etc/inet/ntp.client file, 58
/etc/inet/ntp.conf file, 58
/etc/inet/ntp.drift file, 58
/etc/inet/ntp.keys file, 58
/etc/inet/ntp.server file, 58
/etc/inet/services file, checking for
 UUCP, 581
/etc/inet/slp.conf file, 259, 266, 267, 269,
 270, 271, 273, 274, 275, 276, 278, 280, 281, 283,
 286, 287, 289, 294, 295
/etc/inetd.conf file, 581
/etc/init.d/autofs script, 239
 starting autofs, 163
 stopping autofs, 163
/etc/init.d/ncakmod script, 51
/etc/init.d/ncalogd script, 51
/etc/init.d/nfs.server script
 automatic file-system sharing, 153
 enabling NFS server logging, 156
 enabling WebNFS service, 154
 starting NFS service, 162
 stopping NFS service, 163
/etc/init.d/slpd script, 283
/etc/init.d/slpd script, 267, 269, 270, 271,
 273, 274, 275, 276, 278, 280, 285, 289, 295
/etc/init.d/xntpd script, 58
/etc/mail/aliases.db file, 334, 366
/etc/mail/aliases.dir file, 334, 366
/etc/mail/aliases file, 360, 366, 377, 378
 UUCP and, 582
/etc/mail/aliases.pag file, 334, 367
/etc/mail directory, contents of, 366
/etc/mail/helpfile file, 367, 417
/etc/mail/local-host-names file, 367,
 417
/etc/mail/Mail.rc file, 366
/etc/mail/mailx.rc file, 367
/etc/mail/main.cf file, 367
/etc/mail/relay-domains file, 367
/etc/mail/sendmail.cf file, 367
/etc/mail/sendmail.ct file, 417
/etc/mail/sendmail.cw file, 417
/etc/mail/sendmail.hf file, 417
/etc/mail/sendmail.pid file, 367
/etc/mail/sendmail.st file, 367
/etc/mail/submit.cf file, 367, 390
/etc/mail/subsidiary.cf file, 312, 367,
 376
/etc/mail/trusted-users file, 367, 417
/etc/mnttab file
 comparing with auto_master map, 238
 creating, 216
 description, 198
/etc/nca/nca.if file, 51
/etc/nca/ncakmod.conf file, 51
/etc/nca/ncalogd.conf file, 51
/etc/nca/ncaport.conf file, 51
/etc/netconfig file, description, 198
/etc/nfs/nfslog.conf file
 description, 198
 enabling NFS server logging, 155
 parameters, 199
/etc/nfs/nfslogtab file, description, 198
/etc/nfssec.conf file, description, 198
/etc/nsswitch.conf file, 322, 659
/etc/passwd file, 666
 enabling UUCP logins, 578
/etc/ppp/chap-secrets file
 addressing
 by sppp unit number, 554
 static, 553
 creating
 for a dial-in server, 491
 for trusted callers, 494
 definition, 524
 example, for a PPPoE access server, 562
 syntax, 550
/etc/ppp/myisp-chat.tmpl template, 538
/etc/ppp/options file
 creating
 for a dial-in server, 471
 for a dial-out machine, 463
 definition, 524, 528
/etc/ppp/options.tmpl template, 528
 example PPPoE, 561
 list of examples, 529
 modifying
 for CHAP authentication, 492, 494
 for PAP authentication, 485, 488
 privileges, 526
/etc/ppp/options.tmpl template, 528
/etc/ppp/options.ttya.tmpl
 template, 530

- /etc/ppp/options.*hostname* file
 - creating
 - for a dial-in server, 472
 - for a dial-out machine, 464, 530
 - definition, 524, 529
 - dynamic addressing, 552
 - for a dial-in server, 529
 - list of examples, 531
 - privileges, 526
- /etc/ppp/pap-secrets file
 - addressing
 - by sppp unit number, 554
 - static, 553
 - creating
 - for a dial-in server, 484
 - for a PPPoE access server, 503
 - for trusted callers, 487
 - definition, 524
 - example, for a PPPoE access server, 561
 - syntax, 546
- /etc/ppp/peers directory, 524
- /etc/ppp/peers/myisp.tmpl
 - template, 534
- /etc/ppp/peers/*peer-name* file
 - creating
 - for a dial-out machine, 466
 - for an endpoint on a leased line link, 478
 - definition, 524, 532
 - example, for a PPPoE client, 563
 - list of examples, 535
 - modifying
 - for a PPPoE client, 499
 - for CHAP authentication, 495
 - for PAP authentication, 489
 - privileges, 526
 - useful options, 533
- /etc/ppp/pppoe.*device* file
 - definition, 559
 - for an access server, 503
 - syntax, 559
- /etc/ppp/pppoe file
 - example, 558, 560
 - listing services, 502
 - modifying, 502
 - syntax, 558
- /etc/ppp/pppoe.if file
 - creating
 - for an access server, 501
- creating (Continued)
 - on a PPPoE client, 499
 - definition, 555
 - example, 555
- /etc/project file, 70
- /etc/rmtab file, 198
- /etc/services, nfsd entries, 192
- /etc/shells file, 343, 370
- /etc/syslog.conf file, 347
- /etc/user_attr file, 70
- /etc/uucp/Config file
 - description, 574, 613
 - format, 613
- /etc/uucp/Devconfig file
 - description, 574, 616
 - format, 616
- /etc/uucp/Devices file
 - Class field, 595
 - description, 574, 575, 593
 - Dialer-Token-Pairs field, 596, 598
 - example, for an asppp configuration, 567
 - format, 593
 - Line2 field, 595
 - Line field, 595
 - protocol definitions, 598, 599
 - Systems file Speed field and, 590
 - Systems file Type field and, 594
 - Type field, 594, 595
- /etc/uucp/Dialcodes file
 - description, 575, 603
- /etc/uucp/Dialers file, 599, 602
 - code example, 600
 - description, 575, 599
 - Devices file DTP field and, 597
 - example, for asppp configuration, 567
- /etc/uucp/Grades file, 616
 - default grade, 615
 - description, 575, 614
 - ID-list field, 615, 616
 - Job-size field, 615
 - keywords, 615, 616
 - Permit-type field, 615
 - System-job-grade field, 614, 615
 - User-job-grade field, 614
- /etc/uucp/Limits file
 - description, 575, 617
 - format, 617
- /etc/uucp/Permissions file, 605, 613

/etc/uucp/Permissions file (Continued)

- CALLBACK option, 608, 609
- changing node name, 607
- COMMANDS option, 609, 610, 613
- considerations, 606
- description, 575, 605
- dialback permissions, 608, 609
- file transfer permissions, 606, 608
- format, 605
- forwarding operation, 613
- LOGNAME
 - combining with MACHINE, 612
 - description, 606
 - login IDs for remote computers, 606
- MACHINE
 - combining with LOGNAME, 612
 - default permissions or restrictions, 606
 - description, 606
 - OTHER option, 612
- MYNAME option, 607
- NOREAD option, 608
- NOWRITE option, 608
- OTHER option, 612
- READ option, 608
- remote execution permissions, 609, 612
- REQUEST option, 606
- security set up, 582
- SENDFILES option, 607
- structuring entries, 605
- uucheck command and, 573
- uuxqt daemon and, 573
- VALIDATE option, 610, 612
- WRITE option, 608

/etc/uucp/Poll file

- description, 575, 613
- format, 613

/etc/uucp/Sysfiles file

- description, 575, 604
- format, 604
- printing Systems list, 605
- samples, 604

/etc/uucp/Sysname file, 575, 605

/etc/uucp/Systems file, 587, 593

- Chat Script field, 590, 592
- description, 575, 587
- Devices file Class field and, 595
- Devices file Type field and, 594
- dial-code abbreviations, 575

/etc/uucp/Systems file (Continued)

- escape characters, 591
- example, for an asppp configuration, 566
- format, 588
- hardware flow control, 593
- multiple or different files, 575, 587, 604
- parity setting, 593
- Phone field, 590
- Speed field, 589, 590
- System-Name field, 588
- TCP/IP configuration, 581
- Time field
 - >>>Never entry, 607
 - description, 588
 - Never entry, 588
 - troubleshooting, 584
 - Type field, 589

/etc/vfstab file, 198

- automount command and, 238
- enabling client side failover, 160
- mail clients and, 364
- mounting by diskless clients, 145
- mounting file systems at boot time, 157
- NFS servers and, 158
- nolargefiles option, 160

Ethernet, testing mail configuration on, 345

etrn script, 370

exacct file, 82

example, PPP configurations

- See configuration examples for PPP

executable maps, 247

execute (X.) UUCP files

- cleanup, 580
- description, 619
- uuxqt execution, 573

exit command, 665, 666

expect field of Chat Script field, 590, 591

extended accounting

- activating, 85
- chargeback, 81
- deactivating, 87
- file format, 82
- overview, 81

extended accounting commands, 84

extended accounting status, displaying, 85

F

-F option, unshareall command, 215

f protocol in Devices file, 598

failover

error message, 193

mount command example, 208

NFS support, 147

fair share scheduler

See FSS

FallBackMXhost option, sendmail

command, 395

FastSplit option, sendmail command, 395

FEATURE() declarations

supported, 409

unsupported, 411

feature directory, 369

fg option, mount command, 206

fhtable parameter,

/etc/default/nfslog.conf file, 199

file attributes and NFS version 3, 146

FILE EXISTS message, 620

file permissions

NFS version 3 improvement, 146

WebNFS and, 167

file sharing, 210, 215

examples, 213, 215

giving root access, 212

listed clients only, 211

multiple file systems, 215

NFS version 3 improvements, 146, 147

options, 210

overview, 210

read-only access, 210, 213

read-write access, 210, 213

replicating shared files across several
servers, 179

security issues, 210, 212, 229

unauthenticated users and, 211

unsharing, 214, 215

file system sharing, automatic, 152

file systems

network statistics for, 687, 689

file too large message, 193

file transfer size negotiation, 223

file transfers (UUCP)

daemon, 572

permissions, 606, 608

troubleshooting, 583, 584

file transfers (UUCP) (Continued)

work files C., 618, 619

files and file systems

abbreviated path names, 673, 674

autofs access

NFS file systems using CacheFS, 173, 174

non NFS file systems, 172, 173

autofs selection of files, 242, 244

consolidating project-related files, 176

default file system type, 197

file systems defined, 144

local file systems

default file-system type, 197

unmounting groups, 209

NFS ASCII files and their functions, 198

NFS files and their functions, 197

NFS treatment of, 144

remote file systems

default types, 198

list of remotely mounted file systems, 198

listing clients with remotely mounted file
systems, 216

mounting from file system table, 209

unmounting groups, 209

sharing automatically, 152

find command, searching .rhosts files, 662

firewalls

mounting file systems through, 161

NFS access through, 148

WebNFS access through, 167

flow control hardware

Dialers file, 602

Systems file, 593

forcedirectio option, mount

command, 206

foreground file mounting option, 206

.forward+detail file, 381

.forward files, 379

administering, 341

changing search path, 342

disabling, 342

.forward.hostname files, 380

forwarding operation (UUCP), 613

Fr Time field entry, 588

fslp.conf file, 267

FSS, 101

and processor sets, 107

configuration, 111

FSS (Continued)

- project.cpu-shares, 102
- share definition, 102
- fstypes file, description, 198
- ftp archive, WebNFS and, 167
- ftp command
 - authenticating remote logins, 666
 - interrupting logins, 658
 - opening remote system connections, 667, 668
 - remote logins compared to rlogin and rcp, 666
- FTP Server, nowait, 652
- ftp sessions
 - anonymous ftp accounts, 666
 - closing remote system connections, 668
 - copying files
 - from remote system, 669
 - to remote system, 670
 - opening remote system connections, 668
- ftp sub commands, description, 667
- ftphosts, 637
- fuser command, umountall command and, 209

G

- g option, lockd daemon, 201
- G option, sendmail command, 392
- g protocol in Devices file, 598
- gen-etc-shells script, 343
- generic.m4 file, 368
- generics_entire_domain FEATURE()
 - declaration, 410
- genericstable FEATURE() declaration, 412
- get command (FTP)
 - copying from remote systems, 669
 - example, 669
- gethostbyname command, 385
- Grades file, 614, 616
 - default grade, 615
 - description, 575, 614
 - ID-list field, 615, 616
 - Job-size field, 615
 - keywords, 615, 616
 - Permit-type field, 615
 - System-job-grade field, 614, 615

Grades file (Continued)

- User-job-grade field, 614
- Group keyword of Permit-type field, 616
- GSS-API, and NFS, 148
- guest ftp, setting up, 639

H

- H escape character, 591
- h option
 - umountall command, 209, 210
- hangup, ignoring, 591
- hard option, mount command, 207
- hardware
 - flow control
 - Dialers file, 602
 - Systems file, 593
 - UUCP
 - configurations, 571
 - port selector, 594
- helpfile file, 367
 - sendmail command, 417
- hierarchical mountpoints message, 190
- hierarchical mounts (multiple mounts), 241
- /home directory
 - autofs and, 174
 - NFS server setup, 175
 - /home mount point, 232, 233
- HOST map variable, 245
- host not responding message, 190
- hostname.interface file, NCA and, 50
- hosts
 - checking response of, 682
 - in /etc/hosts.equiv file, 659, 660
 - sending packets to, 682, 683
 - unmounting all file systems from, 209
- hosts.equiv file, 659, 660
- hosts file, 51
- hsfs option, autofs maps, 173
- HTML file, WebNFS and, 167
- httpd command
 - firewall access and WebNFS, 167
 - NCA and, 52
- hung programs, 194
- hyphen (-)
 - dial-code abbreviation, 590
 - Line2 field placeholder, 595

hyphen (-) (Continued)
Speed field placeholder, 590

I

ICMP protocol, 684
ID-list field of Grades file, 615
ID-list field of Grades file, 616
IDLE_TIME parameter,
 /etc/default/nfslogd file, 199
IGMP protocol, 684
ignoring invalid option message, 193
implementing resource pools, 118
in.comsat daemon, 370
in.uucpd daemon, 573
inbound communications
 callback security, 608, 609
 enabling through UUCP chat script, 592
index option
 bad argument error message, 192
 in dfstab file, 154
 WebNFS and, 167
 without public option error message, 192
indirect maps (autofs)
 comments in, 236
 description, 170
 example, 236, 237
 modifying, 171
 overview, 236, 237
 syntax, 236
 when to run automount command, 170
indirect remote logins, 661
inetd.conf file, 581
inetd daemon, in.uucpd invoked by, 573
init command, PPP and, 479
interfaces (PPP)
 asynchronous interface for PPP dial-in, 431
 asynchronous interface for PPP dial-out, 430
 configuring for a PPPoE access server, 501,
 555
 configuring for a PPPoE client
 See also /etc/ppp/pppoe.if file
 HSI/S configuration script, 477
 plumbing PPPoE interfaces with
 /usr/sbin/sppptun, 556
 restricting an interface to PPPoE clients, 502

interfaces (PPP) (Continued)
 synchronous
 for leased lines, 433
interrupting remote logins, 658
-intr option, mount command, 182
IP routing table, 686
IPv6 addresses, sendmail command, 418
ISDN on a PPP link, 431
 chat script for a terminal adapter (TA), 542,
 543

J

Job-size field of Grades file, 615

K

K escape character
 Dialers file, 601
 Systems file chat script, 591
-k option, umountall command, 209
KERB authentication, NFS and, 148
kernel, checking response on server, 183
 /kernel/fs file, checking, 198
keyboard interruption of mounting, 182
keyed map file, creating, 335
keylogin command
 enabling secure NFS, 164
 remote login security issues, 232
keylogout command, secure NFS and, 232
keys file, NTP, 58
keyserv daemon, enabling secure NFS, 164
keywords
 Devices file Type field, 594, 595
 Grades file, 615, 616

L

-l option
 cu command, 583
 nfsd daemon, 202
 umountall command, 209
-L tag option, sendmail command, 392
large files
 disabling creation of, 159

- large files (Continued)
 - NFS support, 147
 - overview, 226
- largefiles option, 206
 - error message, 193
- LCK UUCP lock files, 618
- LDAP, sendmail command and, 415
- ldap_routing FEATURE() declaration, 410
- LDAPDefaultSpec option, sendmail
 - command, 395
- leading space in map entry message, 189
- leased-line link
 - authentication for the link, 435
 - communications process, 434
 - configuring
 - endpoint on the link, 478
 - synchronous interface, 476
 - CSU/DSU, 433
 - definition, 432
 - demand script, 479
 - diagnosing common problems, 509, 510, 511, 512, 518
 - example configuration, 447
 - hardware, 446
 - media, 433
 - parts of the link, 432
 - planning, 446, 447, 449, 478
 - task map for configuring, 475
- legacy services (SLP)
 - advertising, 293, 297
 - definition, 293
- libexacct, 82
- libslp.so library, 256
- Limits file
 - description, 575, 617
 - format, 617
- Line2 field of Devices file, 595
- Line field of Devices file, 595
- link types in PPP
 - comparison of dialup and leased lines, 432
 - dialup, 429
 - leased line, 432
 - parts of a link, 429
 - physical link media, 429
- linking remote logins, 661
- list of tasks, NCA, 44
- listing
 - clients with remotely mounted file systems, 216
 - mounted file systems, 208
 - remotely mounted file systems, 198
 - shared file systems, 213
- local area network (LAN), UUCP
 - configuration, 572
- local cache and NFS version 3, 146
- local delivery agent, mail services, 356
- LOCAL_DOMAIN() m4 configuration
 - macro, 408
- local file systems
 - default file-system type, 197
 - unmounting groups, 209
- local files, updating autofs maps, 170
- local-host-names file, 367, 417
- local_lmtp FEATURE() declaration, 410
- local mail addresses, 361
- local mail alias file, setting up, 333
- local_no_masquerade FEATURE()
 - declaration, 410
- local option (PPP), 479
- lock (LCK) UUCP files, 618
- lockd daemon
 - description, 201
 - options, 201
- LOCKD_GRACE_PERIOD parameter, lockd
 - daemon, 201
- LOCKD_RETRANSMIT_TIMEOUT parameter,
 - lockd daemon, 201
- LOCKD_SERVERS parameter, lockd
 - daemon, 202
- locking, NFS version 3 improvements, 147
- locks, removing, 205
- log file, for NCA, 51
- log levels, sendmail.cf file, 376
- log option
 - in dfstab file, 155
 - share command, 212
- log parameter, /etc/default/nfslog.conf
 - file, 199
- logformat parameter,
 - /etc/default/nfslog.conf file, 199
- logging
 - displaying UUCP log files, 573
 - UUCP log file cleanup, 580

- logging in
 - remote logins, 658, 668
 - authentication (rlogin), 658, 661
 - direct or indirect (rlogin), 661
 - finding who is logged in, 664
 - ftp command, 667
 - interrupting, 658
 - linking logins, 661
 - opening ftp connection, 667, 668
 - using rlogin, 665
- logging out (remote systems), 665, 666
- login command, secure NFS and, 232
- LOGIN FAILED message, 621
- login option (PPP)
 - in /etc/ppp/options for a dial-in server, 486
 - in /etc/ppp/pap-secrets, 489, 549
- logins (UUCP)
 - adding, 578
 - privileged, 611
- LOGNAME Permissions file
 - combining with MACHINE, 612
 - description, 606
 - login IDs for remote computers, 606
 - SENDFILES option, 607
 - VALIDATE option, 610, 612
- lookupdotdomain FEATURE()
 - declaration, 410
- loops, aliases, 346

M

- m4 directory, 369
- M escape character, 591
- m escape character, 592
- MACHINE Permissions file
 - combining with LOGNAME, 612
 - COMMANDS option, 609, 610
 - default permissions or restrictions, 606
 - description, 606
 - OTHER option, 612
- macros
 - defined macros (sendmail), 406
 - for building the configuration file (sendmail), 407
 - m4 configuration macros (sendmail), 408
 - MAX macros (sendmail), 408

- mail addresses
 - % in, 361
 - case sensitivity, 358
 - description, 357
 - domains and subdomains, 358
 - local, 361
 - mail routing and, 382
- mail alias files
 - administering, 326
 - description, 376
 - /etc/mail/aliases file, 377
 - .mailrc aliases, 377
- mail clients
 - definition, 364
 - NFS-mounted file systems and, 317
 - remote mode, 364
 - setting up a mail client, 317
- mail command, 366
- mail commands, interactions of, 370
- mail configuration
 - local mail and a remote connection, 312
 - local only, 312
 - testing, 344
 - typical, 307
- mail connections to other systems, testing, 347
- mail domains
 - name service domains and, 384
 - sendmail.cf file and, 383
- mail exchange (MX) records, 322
- mail gateways
 - configuring, 364
 - definition, 364
 - sendmail.cf file and, 364
 - setting up a mail gateway, 320
 - testing, 345
- mail hosts
 - description, 363
 - setting up a mail host, 319
- mail.local command, 367, 418
- mail queue
 - administering the queue directories, 338
 - displaying the contents of, 339
 - forcing mail queue processing, 339
 - moving the mail queue, 340
 - running a subset of, 339
 - running the old mail queue, 341
- Mail.rc file, 366
- mail routing, mail addresses and, 382

- mail servers, 364
 - backups and, 364
 - description, 363
 - mailboxes on, 361, 364
 - setting up a mail server, 315
 - space requirements, 364
- mail services
 - changes to `mail.local`, 418
 - changes to `mailstats`, 419
 - changes to `makemap`, 419
 - changes to `sendmail`, 389
 - hardware components
 - mail client, 364
 - mail gateway, 364
 - mail host, 363
 - mail server, 363
 - required elements, 362
 - new command, `editmap`, 420
 - planning your mail system, 311
 - software components, 355
 - local delivery agent, 356
 - mail addresses, 357
 - mail aliases, 361
 - mail transfer agent, 355
 - mail user agent, 355
 - mailbox files, 360
 - mailers, 356
 - task maps
 - administering `.forward` files, 341
 - administering mail alias files, 326
 - administering the queue directories, 338
 - comprehensive task map, 310
 - setting up mail services, 314
 - troubleshooting procedures and tips, 344
- mail transfer agents, 355
- mail user agents, 355
- mailbox names
 - owner-owner, 360
 - with `-request` suffix, 360
 - with `owner-` prefix, 360
- MailboxDatabase option, `sendmail` command, 396
- mailboxes
 - files for, 360, 370
 - mail servers and, 364
 - space requirements, 364
- mailcompat filter, 366
- MAILER-DAEMON messages, 348
- MAILER() declarations, 412
- mailer directory, 369
- mailers
 - built-in (`sendmail`)
 - [TCP] and [IPC], 416
 - definition, 356
 - `mail.local` mailer, 367
 - Simple Mail Transfer Protocol (SMTP) mailers, 356
 - Solaris mailers, 356
 - UNIX-to-UNIX Copy command (UUCP) mailers, 357
- `mailq` command, 366
- `.mailrc` aliases, 377
- `.mailrc` file, 362
- `mailstats` command, 366, 419
- `mailtool` command, 370
- `mailx` command, 366
- `mailx.rc` file, 367
- `main.cf` file, 313, 367, 375
- `main.mc` file, 368, 417
- `main-v7sun.mc` file, 417
- maintaining UUCP
 - adding logs, 578
 - mail, 582
 - public directory, 583
 - regular maintenance, 582, 583
 - shell scripts, 579, 581
- `makefile` file, 368
- `makemap` command, 370, 419
- map key bad message, 190
- MAPPING_UPDATE_INTERVAL parameter, `/etc/default/nfslogd` file, 199
- maps (autofs)
 - administrative tasks, 170, 249
 - automount command
 - when to run, 170
 - avoiding mount conflicts, 172
 - comments in, 233, 235, 236
 - default autofs behavior, 247, 249
 - direct, 234, 235
 - executable, 247
 - indirect, 236, 237
 - maintenance methods, 170
 - master, 232, 233
 - modifying, 247
 - direct maps, 171
 - indirect maps, 171

- modifying (Continued)
 - master map, 171
 - multiple mounts, 241
 - network navigation, 240
 - referring to other maps, 245, 246
 - selecting read-only files for clients, 242, 244
 - special characters, 250
 - splitting long lines in, 233, 235, 236
 - starting the navigation process, 233, 240
 - types and their uses, 170
 - variables, 244, 245
- MASQUERADE_EXCEPTION() m4 configuration macro, 408
- master map (auto_master)
 - /- mount point, 232, 236
 - comments in, 233
 - comparing with /etc/mnttab file, 238
 - contents, 232, 234
 - description, 170
 - enabling secure NFS, 165
 - modifying, 171
 - overriding options, 174
 - overview, 232, 233
 - preinstalled, 174
 - security restrictions, 179
 - syntax, 232
 - when to run automount command, 170
- MAX_LOGS_PRESERVE parameter, /etc/default/nfslogd file, 199
- MAXBADCOMMANDS macro, sendmail command, 408
- MAXETRNCOMMANDS macro, sendmail command, 408
- MaxHeadersLength option, sendmail command, 396
- MAXHELOCOMMANDS macro, sendmail command, 408
- MaxMimeHeaderLength option, sendmail command, 396
- MAXNOOPCOMMANDS macro, sendmail command, 408
- MaxQueueChildren option, sendmail command, 396
- MaxRecipientsPerMessage option, sendmail command, 396
- MaxRunnersPerQueue option, sendmail command, 397
- MAXVRFYCOMMANDS macro, sendmail command, 408
- mconnect command, 347, 366
- message types, SLP, 301
- messages
 - UUCP
 - ASSERT error messages, 619, 621
 - checking error messages, 585
 - STATUS error messages, 621, 622
- MeToo option, sendmail command, 400
- mget command (FTP)
 - copying from remote systems, 669
 - example, 670
- MIN_PROCESSINGSIZE parameter, /etc/default/nfslogd file, 199
- minus sign (-), /etc/hosts.equiv file syntax, 659
- mnttab file
 - comparing with auto_master map, 238
 - creating, 216
 - description, 198
- Mo Time field entry, 588
- modem, fixing modem problems, 514
- modem (PPP)
 - chat script examples, 464
 - configuring
 - dial-in server, 468
 - dial-out machine, 462
 - creating chat scripts, 536, 537, 538, 539, 540, 541, 542, 543
 - DSL, 439
 - programming, 462
 - setting the modem speed, 469, 535
- modem (UUCP)
 - direct connection, 597
 - port selector connection, 597, 598
 - setting characteristics, 593, 602
 - troubleshooting, 583
 - UUCP databases
 - DTP field of Devices file, 597, 598
 - UUCP hardware configuration, 572
- modifying
 - direct autofs map, 171
 - indirect autofs map, 171
 - master map (auto_master), 171
- mount command
 - autofs and, 145
 - description, 205

- mount command (Continued)
 - disabling large file creation, 159
 - diskless clients need for, 145
 - failover with, 208
 - manually mounting file systems, 158
 - NFS URL with, 208
 - options
 - file transfer size, 223
 - large files, 206
 - NFS file systems, 206, 207
 - NFS version, 222
 - no arguments, 208
 - nolargefiles, 159, 206
 - public, 161
 - public file handle, 206
 - security mode selection, 207
 - transport protocol, 223
 - using, 207
 - with NFS URL, 162
- mount of server:pathname error, 190
- mount points
 - /- as master map mount point, 232, 236
 - avoiding conflicts, 172
 - /home, 232, 233
 - /net, 234
- mountall command, 209
- mountd daemon
 - checking response on server, 185
 - enabling without rebooting, 187
 - not registered with rpcbind, 193
 - remote mounting requirement, 182
 - verifying if running, 186, 193
- mounting
 - all file systems in a table, 209
 - autofs and, 145, 241
 - background retries, 206
 - diskless client requirements, 145
 - examples, 207, 209
 - force direct I/O, 206
 - foreground retries, 206
 - keyboard interruption during, 182
 - list of mounted file systems, 198
 - nfsd daemon and, 223
 - options for NFS file systems, 206
 - overlaying already mounted file system, 207
 - portmapper and, 223
 - public file handle and, 224
 - read-only specification, 207

- mounting (Continued)
 - read-write specification, 207
 - remote mounting
 - daemons required, 182
 - troubleshooting, 183, 186
 - server not responding, 207
 - soft versus hard, 183
 - /var/mail directory, 317
- mounting file systems
 - autofs and, 159
 - boot time method, 157
 - disabling access for one client, 160
 - manually (on the fly), 158
 - NFS URL with, 161
 - overview, 156
 - task map, 156
 - through a firewall, 161
- mput command (FTP)
 - copying to remote systems, 670
 - example, 672
- mqueue directory, 370
- MS-DOS files, accessing with autofs, 173
- multicast, DAs (SLP), 269
- multicast (SLP), 272, 274, 275, 284, 287, 288
- multihomed hosts (SLP), 276, 285, 287, 288, 291
 - proxy advertisement, 290
 - unicast routing disabled, 289
- multiple files (ftp), 669
- MYNAME option of Permissions file, 607

N

- N escape character
 - Dialers file, 601
- n escape character
 - Dialers file, 601
- N escape character
 - Systems file chat script, 592
- n escape character
 - Systems file chat script, 592
- n option, automountd command, 181
- name option (PPP), 486
- name option (PPP), in /etc/ppp/options for CHAP authentication, 492
- name option (PPP)
 - in /etc/ppp/pap-secrets, 489
 - with noservice, 561

- name service domains, mail domains and, 384
- name services
 - autofs map maintenance methods, 170
 - autofs use of, 248
- name spaces
 - accessing shared, 177
 - autofs and, 149
- names/naming
 - node name
 - UUCP alias, 575, 607
 - UUCP remote computer, 588, 605
- navigating using maps
 - overview, 240
 - starting the process, 233, 240
- NCA
 - architecture, 52
 - changing logging, 49
 - disabling, 48
 - enabling, 46
 - files description, 50
 - httpd and, 52
 - kernel module, 52
 - list of tasks, 44
 - new features, 44
 - overview, 43
 - requirements, 45
 - socket library, 49
 - sockets, 45
- nca_addr.so library, 51
- nca_httpd_1.door file, 51
- nca.if file, 46, 51
- NCA log file, 51
- ncab2clf command, 51
- ncaconfd command, 51
- ncakmod.conf file, 46, 49, 51
- ncakmod module, 52
- ncalogd.conf file, 47, 49, 51
- ncalogd script, 51
- ncaport.conf file, 51
- negotiation
 - file transfer size, 223
 - NFS version, 222
 - transport protocol, 222
 - WebNFS security, 149
- /net mount point
 - access method, 234
 - description, 234
- net.slp.DAActiveDiscoveryInterval
 - property, 269
 - definition, 268
- net.slp.DAAddresses property, 272, 282, 287
 - definition, 268
- net.slp.DAAttributes property, 273
- net.slp.DAHeartBeat property, 271, 273
 - definition, 268
- net.slp.interfaces property, 286, 290, 291
 - configuring, 288
- net.slp.isBroadcastOnly property, 276, 288
- net.slp.isDA property, 267
- net.slp.MTU property, 275
- net.slp.multicastTTL property, 274
- net.slp.passiveDADetection property, 269
 - definition, 268
- net.slp.randomWaitBound property, 279
- net.slp.serializedRegURL property, 295
- net.slp.useScopes property, 282, 296
 - definition, 281
- netconfig file, description, 198
- netstat command, 261, 683, 686
 - i option (interfaces), 683, 684
 - overview, 681, 683
 - r option (IP routing table), 686
 - s option (per protocol), 684
- network authentication for remote logins, 658, 660, 661
- Network Cache and Accelerator
 - See NCA
- network databases
 - services
 - UUCP port, 581
- network interfaces (SLP), nonrouted
 - considerations, 291
- network lock manager, 147
- networks
 - commands for monitoring performance, 681
 - displaying performance information, 681, 682, 683, 689
 - client statistics, 687, 689
 - collision rate, 684
 - host response, 682
 - interface statistics, 683, 686
 - IP routing table, 686
 - server statistics, 687, 689
 - packets
 - capturing from network, 681, 683

- packets (Continued)
 - dropped, 683
 - error rates, 684
 - number transmitted, 684
 - reliability testing, 681, 682, 683
 - sending to hosts, 682, 683
- tracing client calls to servers, 681, 683
- troubleshooting
 - hardware components, 689
 - high retransmission rate, 686
- Never Time field entry, 588, 607
- newaliases command, UUCP and, 582
- newaliases link, 370
- newkey command, enabling secure NFS, 164
- newline escape characters, 592, 601
- NFS ACL, 147
- NFS administration, administrator
 - responsibilities, 152
- NFS can't support nolargefiles
 - message, 193
- NFS client
 - disabling
 - autofs browsability, 180
- nfs.client, lockd and, 201
- NFS clients
 - incompatible operating system support, 178
 - NFS services, 143
- NFS environment
 - Secure NFS system, 229
- nfs file
 - See /etc/default/nfs file
- NFS locking, client-side failover and, 226
- NFS-mounted file systems
 - mail clients and, 315, 317
 - mail servers and, 315
- NFS server logging
 - enabling, 155
 - overview, 149
- nfs.server script
 - automatic file-system sharing, 153
 - enabling NFS server logging, 156
 - enabling WebNFS service, 154
 - nfsd and, 202
 - starting NFS service, 162
 - stopping NFS service, 163
- NFS servers
 - autofs selection of files, 244
 - daemons required for remote mounting, 182

- NFS servers (Continued)
 - identifying current, 188
 - maintaining, 152
 - not responding during mounting, 207
 - replicating shared files, 179
 - troubleshooting
 - clearing problems, 183
 - remote mounting problems, 183, 193
 - weighting in maps, 244
- NFS service
 - benefits, 144
 - features, 146
 - file systems, 144
 - overview, 144
 - servers and clients, 143
 - version 2 protocol, 146
 - version 3 protocol, 146
- NFS services
 - restarting, 187
 - starting, 162
 - stopping, 162
 - task map, 162
- NFS troubleshooting
 - determining where NFS service has
 - failed, 186
 - hung programs, 194
 - remote mounting problems, 193
 - server problems, 183
 - strategies, 182
- NFS URL
 - autofs and, 180
 - mount command example, 208
 - mounting file systems with, 161
 - mounting with, 149
 - syntax, 167
 - WebNFS and, 166
- NFS V2 can't support largefiles
 - message, 193
- NFS version negotiation, 222
- nfscast: cannot receive reply message, 191
- nfscast: cannot send packet message, 191
- nfscast: select message, 191
- nfsd daemon
 - checking response on server, 184
 - description, 202
 - enabling without rebooting, 187
 - mounting and, 223
 - options, 202

- nfsd daemon (Continued)
 - remote mounting requirement, 182
 - verifying if running, 186
- NFSD_LISTEN_BACKLOG parameter, nfsd daemon, 202
- NFSD_MAX_CONNECTIONS parameter, nfsd daemon, 202
- NFSD_SERVERS parameter, nfsd daemon, 202
- nfslog.conf file
 - description, 198
 - enabling NFS server logging, 155
 - parameters, 199
- nfslogd daemon
 - description, 203
 - enabling NFS server logging, 156
- nfslogd file, description, 197
- nfslogtab file, description, 198
- nfssec.conf file, description, 198
- nfsstat command, 188, 217, 687, 689
 - c option (clients), 687
 - m option (per file system), 687, 689
 - overview, 681, 687
 - s option (servers), 687
- NiceQueueRun option, sendmail command, 397
- NIS+ mail_aliases table, 379
 - adding aliases to, 329
 - adding entries by editing, 330
 - deleting entries from, 332
 - editing entries in, 331
 - initiating tables, 328
 - listing an individual entry from, 328
 - listing partial matches, 329
 - listing the entire contents of, 328
 - managing alias entries in, 327
- NIS+ name service, updating autofs maps, 170
- NIS aliases map, 378
- NIS mail_aliases map, setting up, 332
- NIS name service, updating autofs maps, 170
- nisaddcred command, enabling secure NFS, 164
- nistbladm command
 - modifying autofs master map (auto_master), 171
 - modifying direct autofs map, 171
 - modifying indirect autofs map, 171
- nnn escape character, 601
- no_default_msa FEATURE()
 - declaration, 411
- NO DEVICES AVAILABLE message, 621
- no info message, 191
- No such file or directory message, 193
- NO UUCP SERVICE NUMBER message, 620
- noauth option (PPP), 466, 479, 533, 564
- nocanonicalize FEATURE() declaration, 410
- noccp option (PPP), 471, 564
- node name
 - UUCP alias, 575, 607
 - UUCP remote computer, 588, 605
- nodefaultroute option (PPP), 471
- noipdefault option (PPP), 466, 533
- nolargefiles option
 - error message, 193
 - in vfstab file, 160
 - mount command, 159, 206
- Non-group keyword of Permit-type field, 616
- Non-user keyword of Permit-type field, 616
- NOREAD option of Permissions file, 608
- noservice option (PPP), 561
- nosuid option, share command, 212
- Not a directory message, 191
- Not found message, 190
- nouucp FEATURE() declaration, 411
- NOWRITE option of Permissions file, 608
- nservers option, nfsd daemon, 202
- nsswitch.conf file, 322, 659
- nthreads option, lockd daemon, 202
- NTP client, setting up, 56
- ntp.conf file
 - for client, 56
 - for server, 56
- NTP files, 58
- NTP Server, setting up, 56
- ntpdate command, 58
- ntpq command, 58
- ntpstats directory, 58
- ntptrace command, 58
- null escape character, 592, 601
- nullclient FEATURE() declaration, 411
- number sign (#)
 - comments in direct maps, 235
 - comments in indirect maps, 236
 - comments in master map (auto_master), 233

O

- o option
 - mount command, 206, 207
 - share command, 210, 213
- O option of mount command, 207
- octal numbers escape character, 592, 601
- OK message, 621
- open errors, NFS and, 146
- opening remote system connections, 667, 668
- operating systems
 - map variables, 245
 - supporting incompatible versions, 178
- options (PPP)
 - asynmap, 530
 - auth, 486
 - call, 473, 533
 - connect, 466, 533, 543, 563
 - crtstcts, 464
 - defaultroute, 533, 564
 - guidelines for use, 523
 - init, 479, 530
 - local, 479
 - login, 486, 549
 - name, 486, 489
 - noauth, 466, 479, 533, 564
 - noccp, 471, 564
 - nodefaultroute, 471
 - noipdefault, 466, 533
 - noservice, 561
 - option privileges, 526, 527
 - parsing, by the pppd daemon, 525
 - passive, 479
 - persist, 479
 - plugin, 563
 - remotename, 486, 533
 - require-chap, 493
 - require-pap, 486
 - sppptun, 563
 - sync, 479
 - updetach, 534
 - user, 486, 533
 - xonxoff, 472
- options file, in PPP
 - See /etc/ppp/options
- options for PPP, debug, 508
- options in sendmail command
 - ClientPortOptions option, 400
 - command-line options, 391

- options in sendmail command (Continued)
 - configuration file options
 - deprecated and unsupported, 400
 - supported, 392
 - DaemonPortOptions option, 401
 - PidFile option, 403
 - PrivacyOptions option, 403
 - ProcessTitlePrefix option, 403
 - Timeout option, 404
- options.ttyname file (PPP)
 - See /etc/ppp/options.ttyname
- OSNAME map variable, 245
- OSREL map variable, 245
- ostype directory, 369
- OSVERS map variable, 245
- OTHER option of Permissions file, 612
- overlying already mounted file system, 207
- owner- prefix, mail alias with, 361

P

- p escape character
 - Dialers file, 601
 - Systems file chat script, 592
- packet size, configuring for SLP, 275
- PAM (pluggable authentication module),
 - identity management, 71
- PAP credentials database
 - creating
 - for a dial-in server, 483, 484
 - for trusted callers, 487
- parity
 - Dialers file, 602
 - Systems file, 593
- passive mode, 607
- passive option (PPP), 479
- passwd file, enabling UUCP logins, 578
- Password Authentication Protocol (PAP)
 - authentication process, 547
 - configuring
 - on a dial-in server, 485
 - trusted callers, 487, 488, 489
 - creating a PAP credentials database, 483
 - definition, 546
 - /etc/ppp/pap-secrets file, 546
 - example configuration, 450
 - planning, 482

Password Authentication Protocol (PAP)
(Continued)

- suggestions for passwords, 546
- task maps, 482, 483
- using the login option, 549

passwords

- authentication for remote logins
 - ftp command, 666, 668
 - rlogin command, 658, 661, 665
- autofs and superuser passwords, 145
- DH password protection, 229
- Secure RPC password creation, 164
- UUCP privileged, 611

path names

- rcp command
 - absolute or abbreviated, 673, 674
 - syntax options, 674
- tilde (~) in, 673, 674

pathconf: no info message, 191

pathconf: server not responding message, 191

PC-DOS files, accessing with autofs, 173

pcfs option, autofs maps, 173

peer

- access server, 437, 455
- authenticatee, 435
- authenticator, 435
- definition, 429
- dial-in server, 429
- dial-out machine, 429
- leased-line peer, 433
- PPPoE client, 437, 454

peers file for PPP

See /etc/ppp/peers/peer-name

penril entry in Dialers file, 601, 602

percent sign (%) in mailbox names, 361

Permission denied message, 193

permissions

- copying requirements, 675
- NFS version 3 improvement, 146

Permissions file, 605, 613

- CALLBACK option, 608, 609
- changing node name, 607
- COMMANDS option, 609, 610, 613
- considerations, 606
- description, 575, 605
- dialback permissions, 608, 609
- file transfer permissions, 606, 608
- format, 605

Permissions file (Continued)

forwarding operation, 613

LOGNAME

- combining with MACHINE, 612
- description, 606
- login IDs for remote computers, 606

MACHINE

- combining with LOGNAME, 612
- default permissions or restrictions, 606
- description, 606
- OTHER option, 612

MYNAME option, 607

NOREAD option, 608

NOWRITE option, 608

OTHER option, 612

READ option, 608

remote execution permissions, 609, 612

REQUEST option, 606

security set up, 582

SENDFILES option, 607

structuring entries, 605

uucheck command and, 573

uuxqt daemon and, 573

VALIDATE option, 610, 612

WRITE option, 608

Permit-type field of Grades file, 615

persist option (PPP), 479

Phone field of Systems file, 590

PidFile option

sendmail command, 397, 403

ping command, 277, 663, 681, 682

PKCGET READ message, 620

PKXSTART message, 621

pluggable authentication module

See PAM

plugin option (PPP), 563

plus sign (+)

/etc/hosts.equiv file syntax, 659, 660

in autofs map names, 245, 246

Point-to-Point Protocol

See PPP

Poll file

description, 575, 613

format, 613

polling remote computers (UUCP), 575, 613

Port Selector variable in Devices file, 594

portmapper, mounting and, 223

- ports
 - Devices file entry, 595
 - UUCP, 581
- postmaster alias, creating, 335
- postmaster mailbox
 - creating, 336
 - description, 360
 - testing, 345
- pound sign (#)
 - comments in direct maps, 235
 - comments in indirect maps, 236
 - comments in master map (auto_master), 233
- PPP
 - authentication, 434, 435
 - chat script examples, 464
 - common problems, 506
 - compatibility, 426
 - converting from asynchronous PPP, 568
 - dial-up link, 429
 - difference with asppp, 426
 - DSL support, 436
 - file privileges, 525, 527
 - ISDN support, 431
 - leased-line link, 432
 - modem settings, basic, 462
 - options for configuration files
 - See* options (PPP)
 - overview, 425
 - parts of a link, 428
 - pppd
 - See also* pppd command
 - PPPoE, 437
 - problem solving
 - See also* troubleshooting PPP
 - related RFCs, 428
 - resources, external, 427
 - summary of configuration files, 524
 - task map for PPP planning, 441
- pppd command
 - definition, 524
 - initiating a call, 473
 - obtaining diagnostics, 507, 508, 520
 - parsing options, 525
- pppd command, testing a DSL line, 500
- pppd command
 - turning on debugging, 508
- pppdebug log file, 519
- PPPoE
 - configuring an access server, 501, 502, 503, 504
 - DSLAM, 439
 - fixing common problems, 519, 520, 521
 - list of commands and files, 555
 - obtaining snoop traces, 521
 - overview, 437
 - planning for the tunnel, 454, 455, 456, 457
 - providing services from an access server, 502, 558, 559
 - task maps for configuring, 497
 - tunnel participants, 437
- PPPoE client
 - configuring, 499, 500, 562, 563, 564
 - definition, 437
 - equipment, 454
 - /etc/ppp/peers/peer-name file usage (PPPoE), 563
 - planning, 454, 498
 - task map for configuring, 497
- pppoe.so shared object, 560
- pppoe.so shared object, 562
- pppoec utility
 - definition, 562
 - obtaining diagnostics, 520
- pppoed daemon
 - definition, 557
 - starting, 502
- .ppprc file
 - creating, 470
 - definition, 524
 - for a dial-in server, 531
 - for a dial-out machine, 532
 - privileges, 526
- praliases command, 366
- preserve_local_plus_detail FEATURE()
 - declaration, 411
- preserve_luser_host FEATURE()
 - declaration, 411
- printing
 - list of remotely mounted directories, 216
 - list of shared or exported files, 216
- PrivacyOptions option
 - sendmail command, 397, 403
- privilege levels, 91
- processor type map variable, 245

ProcessTitlePrefix option
 sendmail command, 397, 403
 programs, hung, 194
 project
 active state, 102
 definition, 70
 idle state, 102
 with zero shares, 102
 project 0, 107
 project.cpu-shares, 106
 project entry format, 71
 project system
 See project 0
 projects, consolidating files, 176
 protocol definitions in Devices file, 598, 599
 proxy advertisement (SLP), 294, 295
 proxy registration (SLP), 294, 296
 multihomed hosts, 290
 PRUNE_TIMEOUT parameter,
 /etc/default/nfslogd file, 199
 pstack command, 218
 public directory maintenance (UUCP), 583
 public file handle
 autofs and, 180
 mounting and, 224
 NFS mounting with, 148
 WebNFS and, 166
 public-key cryptography
 common key, 231
 conversation key, 231
 database of public keys, 229, 230
 DH authentication, 230, 231
 secret key
 database, 230
 deleting from remote server, 231
 time synchronization, 231
 public option
 in dfstab file, 154
 mount command, 161, 206
 share error message, 194
 WebNFS and, 167
 publickey map, 230
 enabling secure NFS, 164
 put command (FTP)
 copying to remote systems, 670
 example, 671
 putacct, 83

Q

-q[!] *Isubstring* option, sendmail
 command, 392
 -q[!] *Rsubstring* option, sendmail
 command, 392
 -q[!] *Ssubstring* option, sendmail
 command, 392
 -q option, uustat command, 583
 -qf option, sendmail command, 392
 -qGname option, sendmail command, 392
 -qptime option, sendmail command, 392
 queue (UUCP)
 administrative files, 618, 619
 clean-up command, 573
 job grade definitions, 614, 616
 scheduling daemon, 573
 spool directory, 618
 uusched daemon
 description, 573
 maximum simultaneous executions, 575,
 617
 queue features, sendmail command, 414
 QueueFileMode option, sendmail
 command, 397
 queuegroup FEATURE() declaration, 411
 QueueLA option, sendmail command, 397
 QueueSortOrder option, sendmail
 command, 398

R

r escape character
 Dialers file, 601
 Systems file chat script, 592
 r option, mount command, 207
 -r option
 umountall command, 209, 210
 uucp command, 584
 Uutry command, 583
 rbl FEATURE() declaration, 412
 rcp command, 673, 677
 copying between local and remote
 systems, 675, 677
 copying directories, 675
 description, 673
 examples, 676, 677

- rcp command (Continued)
 - path names
 - absolute or abbreviated, 673, 674
 - syntax options, 674
 - security issues, 673
 - specifying source and target, 673, 674
- rctls
 - See* resource controls
- rdate command, 57
- read-only type
 - file selection by autofs, 242, 244
 - mounting file systems as, 207
 - sharing file systems as, 210, 213
- READ option of Permissions file, 608
 - NOREAD option, 608
- read-write type
 - mounting file systems as, 207
 - sharing file systems as, 210, 213
- real ftp, setting up, 638
- recovery from bad pooladm configuration, 123
- RefuseLA option, sendmail command, 398
- registration lifetime (SLP), 261
- relay_mail_from FEATURE()
 - declaration, 411
- relay-domains file, 367
- remote copying
 - using rcp, 673, 677
- REMOTE DOES NOT KNOW ME
 - message, 622
- remote execution (UUCP)
 - commands, 606, 609, 612
 - daemon, 573
 - work files C., 618, 619
- remote file copying, using ftp, 667
- remote file systems
 - default types, 198
 - list of remotely mounted file systems, 198
 - listing clients with remotely mounted file systems, 216
 - unmounting groups, 209
- REMOTE HAS A LCK FILE FOR ME
 - message, 622
- remote logins
 - authentication (ftp), 666
 - authentication (rlogin), 658, 661
 - /etc/hosts.equiv file, 659, 660
 - network authentication or remote system authentication, 658, 659
- authentication (rlogin) (Continued)
 - .rhosts files, 660, 661
 - closing ftp connection, 668
 - direct or indirect (rlogin), 661
 - domains, 658
 - finding who is logged in, 664
 - ftp commands, 667
 - interrupting, 658
 - linking logins, 661
 - opening ftp connection, 667, 668
 - removing .rhosts files, 662
 - using rlogin command, 665
 - verifying remote system operation, 663
- remote_mode FEATURE() declaration, 412
- remote mounting
 - daemons required, 182
 - troubleshooting, 183, 186
- REMOTE REJECT, UNKNOWN MESSAGE
 - message, 622
- REMOTE REJECT AFTER LOGIN
 - message, 622
- remote systems
 - definition, 627
 - logging in, 658, 668
 - logging out (exit), 665, 666
 - overview, 627
 - remote copying
 - using rcp, 673, 677
 - remote file copying
 - using ftp command, 667
 - verifying operation, 663
- remote.unknown file, 617
- remotename option (PPP), 486, 533
- remount message, 190
- removing locks, 205
- removing resource pools, 123
- replicas must have the same version, 194
- replicated file system, 225
- replicated mounts, soft option and, 194
- replicated mounts must be read-only, 194
- replicated mounts must not be soft, 194
- replicating shared files across several servers, 179
- REQUEST option of Permissions file, 606
- Requests for Comments (RFCs), PPP, 428
- require-chap option (PPP), 493
- require-pap option (PPP), 486

- ResolverOptions option, sendmail
 - command, 398
 - resource controls
 - available, 90
 - changing temporarily, 95
 - configuring, 90
 - definition, 89
 - threshold values, 92
 - resource limits, 89
 - resource management, 61
 - constraints, 63
 - partitioning, 64
 - scheduling, 64
 - resource pools, 115
 - administering, 117
 - binding to, 124
 - creating configurations, 119
 - implementing, 118
 - removing, 123
 - resources, shared, 198
 - retry subfield of Time field, 589
 - return escape character, 601
 - RETURN FROM fixline ioctl message, 620
 - .rhosts files
 - deleting, 662
 - description, 660
 - remote system authentication process, 659, 660
 - searching, 662
 - security issues, 660, 661
 - rlimits
 - See resource limits
 - rlogin command
 - authentication, 658, 661
 - /etc/hosts.equiv file, 659, 660
 - network or remote system authentication, 658, 659
 - .rhosts files, 660, 661
 - description, 658
 - direct or indirect logins, 661
 - interrupting logins, 658
 - process after logging in, 662
 - secure NFS and, 232
 - using, 665
 - rm command, 661
 - rmail command, 366
 - rmtab file, description, 198
 - ro option
 - mount command with -o flag, 207
 - share command, 210, 213
 - root directory, mounting by diskless clients, 145
 - root option, share command, 212
 - RPC, 687
 - authentication, 230
 - Secure
 - DH authorization issues, 231, 232
 - overview, 229, 230
 - rpcbind daemon
 - dead or hung, 193
 - mountd daemon not registered, 193
 - warm start, 187
 - rpcinfo command, 219
 - RPCSEC_GSS, 148
 - RrtImpliesDsn option, sendmail
 - command, 398
 - RS-232 telephone lines, UUCP
 - configuration, 571
 - rule sets
 - new, 416
 - testing, 346
 - rusers command, 664
 - rw=client option, share command, 211
 - rw option
 - mount command with -o flag, 207
 - share command, 210
 - share command, 213
- ## S
- s escape character
 - Dialers file, 601
 - Systems file chat script, 592
 - s option, umountall command, 209
 - SA server (SLP), 279
 - Sa Time field entry, 588
 - SAs (SLP), 282, 289, 294
 - scheduling daemon for UUCP, 573
 - scopes (SLP), 271, 281, 282, 284, 291, 294, 296
 - definition, 253
 - deploying, 281
 - scripts
 - chat scripts (UUCP), 592
 - basic script, 590

- chat scripts (UUCP) (Continued)
 - enabling dialback, 592
 - escape characters, 591
 - expect field, 590, 591
 - format, 590
 - shell scripts (UUCP), 579, 581
- searching
 - .rhosts files, 662
 - users logged in to remote system, 664
- sec=dh option
 - auto_master map, 165
 - dfstab file, 165
- secret key
 - database, 230
 - deleting from remote server, 231
 - server crash and, 231
- secrets file for PPP
 - See /etc/ppp/pap-secrets file
 - See also /etc/ppp/chap-secrets file
- secure mounting, dfstab file option, 165
- Secure NFS system
 - administering, 163
- secure NFS system, DH authentication and, 163
- Secure NFS system
 - domain name, 163
 - overview, 229
 - setting up, 164
- Secure RPC
 - DH authorization issues, 231, 232
 - overview, 229, 230
- security
 - applying autofs restrictions, 179
 - copy operation issues, 673
 - DH authentication
 - dfstab file option, 165
 - overview, 230, 231
 - password protection, 229
 - user authentication, 229
 - /etc/hosts.equiv file issues, 660
 - file-sharing issues, 210, 212
 - mount command and, 207
 - NFS and, 147
 - NFS version 3 and, 146
 - .rhosts file issues, 660, 661, 662
 - Secure NFS system
 - administering, 163
 - overview, 229
- security (Continued)
 - Secure RPC
 - DH authorization issues, 231, 232
 - overview, 229, 230
 - UNIX authentication, 229, 230
 - UUCP
 - COMMANDS option of Permissions file, 609, 610
 - setting up, 582
 - sticky bit for public directory files, 583
 - VALIDATE option of Permissions file, 610, 612
- security flavors, 148
- security services, list of, 198
- SENDFILES option of Permissions file, 607
- sendmail.cf file, 367
 - alternate configuration for, 325
 - building the configuration file, 323
 - description, 375
 - log levels, 376
 - mail domains and, 383
 - mail gateways and, 364
 - mail hosts and, 376
 - mail servers and, 376
 - mailers, description, 356
 - vendor setting, 354
 - version level, 354
- sendmail command
 - alternative commands, 353
 - changes in file name or file location, 417
 - changes to, 389
 - command-line options, 391
 - compilation flags, 352
 - configuration file options
 - deprecated and unsupported, 400
 - supported, 392
 - delivery agent flags, 412
 - description, 371
 - equates for delivery agents, 413
 - error messages, 348
 - /etc/mail/helpfile file, 417
 - /etc/mail/local-host-names file, 417
 - /etc/mail/sendmail.ct file, 417
 - /etc/mail/sendmail.cw file, 417
 - /etc/mail/submit.cf, 390
 - /etc/mail/trusted-users file, 417
 - FEATURE() declarations
 - supported, 409

- FEATURE() declarations (Continued)
 - unsupported, 411
 - FEATURE() declarations, changes to, 409
 - features of, 374
 - .forward files, 379
 - helpfile file, 417
 - interactions of NIS+ and, 386
 - interactions of NIS and, 385
 - interactions with NIS+ and DNS, 387
 - interactions with NIS and DNS, 386
 - interactions with other mail commands, 375
 - IPv6 addresses, 418
 - LDAP and, 415
 - local-host-names file, 417
 - macros
 - defined macros, 406
 - for building the configuration file, 407
 - m4 configuration macros, 408
 - MAX macros, 408
 - MAILER() declarations, 412
 - mailers, built-in
 - [TCP] and [IPC], 416
 - main.mc file, 417
 - main-v7sun.mc file, 417
 - name services and, 384
 - NIS+ mail_aliases table, 379
 - NIS aliases map, 378
 - other changes, 421
 - queue features, 414
 - rule sets, 416
 - sendmail.ct file, 417
 - sendmail.cw file, 417
 - submit.cf file, 390
 - subsidiary.mc file, 418
 - subsidiary-v7sun.mc file, 418
 - trusted-users file, 417
 - /usr/lib/mail/cf/main.mc file, 417
 - /usr/lib/mail/cf/main-v7sun.mc file, 417
 - /usr/lib/mail/cf/subsidiary.mc file, 418
 - /usr/lib/mail/cf/subsidiary-v7sun.mc file, 418
- sendmail.ct file, 417
 - sendmail.cw file, 417
 - sendmail.hf file, 417
 - sendmail.pid file, 367, 370
 - sendmail program, new features, 305
 - sendmail.st file, 367
 - sendmailvars.org_dir table, 370
 - SendMimeErrors option, sendmail command, 399
 - serial port
 - configuring
 - dial-out machine, 462
 - for a dial-in server, 468
 - on a dial-in server, 529
 - serial unmounting, 209
 - server consolidation, 65
 - server not responding message, 190, 191
 - hung programs, 194
 - keyboard interrupt for, 182
 - remote mounting problems, 193
 - servers
 - See NFS servers
 - autofs selection of files, 242
 - crashes and secret keys, 231
 - displaying information about, 681, 687, 689
 - home directory server setup, 175
 - NFS servers and vfstab file, 158
 - NFS services, 143
 - tracing client calls to, 681, 683
 - service advertisement (SLP), 273, 295
 - service agent (SLP), 268, 273
 - service discovery (SLP), 276, 278, 284
 - service requests (SLP), 284
 - service URLs
 - proxy registration (SLP), 294, 296
 - services database, UUCP port, 581
 - setgid mode, share command option preventing, 212
 - setmnt command, 216
 - setting up
 - local mail alias file, 333
 - mail client, 317
 - mail gateway, 320
 - mail host, 319
 - mail server, 315
 - NIS mail_aliases map, 332
 - virtual host, 322
 - setuid mode
 - Secure RPC and, 232
 - share command option preventing, 212
 - sh directory, 369
 - share command, 210, 213
 - description, 210

- share command (Continued)
 - options, 210
 - security issues, 210, 212
 - using, 213
- shareall command
 - automatic file-system sharing, 153
 - disabling mount access for one client, 161
 - enabling NFS server logging, 156
 - enabling WebNFS service, 155
 - overview, 215
- shared resources, list of, 198
- SharedMemoryKey option, sendmail
 - command, 399
- sharetab file, description, 198
- shell scripts (UUCP), 579, 581
 - automatic execution, 579
 - running manually, 579
 - uudemon.admin, 580
 - uudemon.cleanup, 580
 - uudemon.hour
 - description, 580
 - uusched daemon execution by, 573
 - uuxqt daemon execution by, 573
 - uudemon.poll, 580, 613
- shells file, 370
- showmount command, 216
- single-user mode and security, 231
- slash (/)
 - /- as master map mount point, 232, 236
 - master map names preceded by, 233
 - root directory
 - mounting by diskless clients, 145
- SLP
 - advertising, 285
 - agents and processes, 254
 - analyzing snoop slp trace, 261
 - architecture, 253
 - configuring, 259, 265, 266, 275, 276
 - daemon, 256
 - discovery requests, 277
 - enabling, 263
 - implementation, 256
 - logging, 253
 - performance tuning, 272
 - planning deployment, 259
- slp.conf file, 266, 272
- slp.jar library, 256
- SLP message types, 301
- SLP status codes, 299
- slpd.conf file, 268, 282
- slpd daemon, 267, 268, 271, 282, 287, 288, 290, 294, 297
 - DAs, 279
 - removing DAs, 271
 - SA server, 279
- SLPv2, interoperability with SLPv1, 285
- SMART_HOST() m4 configuration macro, 409
- smrsh command, 368
- SMTP (Simple Mail Transfer Protocol)
 - mailers, 356
 - sendmail.cf file, 391
- snoop command, 221, 681, 683
 - using with SLP, 260, 261, 263, 273, 279, 287, 289
- snoop trace, for PPPoE, 521
- sockets, NCA and, 45
- soft option, mount command, 207
- Solaris, UUCP version, 587
- solaris2.m4 file, 369
- solaris8.m4 file, 369
- solaris2.ml.m4 file, 369
- solaris2.pre5.m4 file, 369
- solaris-antispam.m4 file, 369
- solaris-generic.m4 file, 342, 343, 369
- Solaris Management Console
 - performance monitoring, 134
 - setting resource controls, 140
- Solaris PPP 4.0
 - See PPP
- space escape character, 592, 601
- special characters in maps, 250
- Speed field
 - Devices file Class field and, 595
 - Systems file, 589, 590
- spool (UUCP)
 - administrative files, 618, 619
 - clean-up command, 573
 - directory, 618
 - job grade definitions, 614, 616
 - uusched daemon
 - description, 573
 - maximum simultaneous executions, 575, 617
- sppp unit number, PPP address assignment, 554
- sppptun option (PPP), 563
- spray command, 681, 682, 683

- starting
 - autofs service, 163
 - enabling dialback through chat script, 592
 - NFS services, 162
 - turning on
 - CLOCAL flag, 591
 - echo checking, 591, 601
 - UUCP shell scripts, 579, 581
 - STARTUP FAILED message, 622
 - statd daemon, 203
 - static addressing, PPP, 553
 - status codes, SLP, 299
 - .Status directory, 585
 - STATUS error messages (UUCP), 585, 621, 622
 - sticky bit for public directory files, 583
 - stopping
 - autofs service, 163
 - NFS services, 162
 - turning off
 - CLOCAL flag, 592
 - echo checking, 591, 601
 - STREAMS
 - device configuration, 616
 - dialer token pairs, 596
 - STTY flow control, 593, 602
 - Su Time field entry, 588
 - submit.cf file, 367, 390
 - submit.mc file, 368
 - subsidiary.cf file, 312, 367, 376
 - subsidiary.mc file, 368, 418
 - subsidiary-v7sun.mc file, 418
 - sun_reverse_alias_files FEATURE()
 - declaration, 412
 - sun_reverse_alias_nis FEATURE()
 - declaration, 412
 - sun_reverse_alias_nisplus FEATURE()
 - declaration, 412
 - SuperSafe option, sendmail command, 399
 - superusers, autofs and passwords, 145
 - sync option (PPP), 479
 - synchronizing time, 231
 - with another system, 57
 - synchronous PPP, *See* leased-line link, 432
 - configuring synchronous devices, 476
 - Sys-Name variable of Type field, 594
 - Sysfiles file
 - description, 575, 604
 - format, 604

- Sysfiles file (Continued)
 - printing Systems list, 605
 - samples, 604
- syslog.conf file, 347
- syslogd command, 370
- SYSLST OVERFLOW message, 620
- Sysname file, 575, 605
- system authentication for remote logins, 658, 659
- System-job-grade field of Grades file, 614, 615
- System-Name field of Systems file, 588
- SYSTEM NOT IN Systems FILE message, 622
- Systems file, 587, 593
 - Chat Script field, 590, 592
 - description, 575, 587
 - Devices file Class field and, 595
 - Devices file Type field and, 594
 - dial-code abbreviations, 575, 590
 - escape characters, 591
 - format, 588
 - hardware flow control, 593
 - multiple or different files, 575, 587, 604
 - parity setting, 593
 - Phone field, 590
 - Speed field, 589, 590
 - System-Name field, 588
 - TCP/IP configuration, 581
 - Time field
 - description, 588
 - Never entry, 588, 607
 - troubleshooting, 584
 - Type field, 589

T

- t escape character, 592
- T escape character
 - Devices file, 598
 - Dialers file, 598, 601
- t option, lockd daemon, 201
- t protocol in Devices file, 598
- tab escape character, 592
- TALKING message, 621
- tasks, resource management, 74
- TCP, NFS version 3 and, 147
- TCP dialer type, 596

- TCP/IP networks
 - UUCP over, 581
- TCP/IP traffic, 681, 683, 684
- TCP protocol, 684
- telephone lines, UUCP configuration, 572
- telephone numbers in `Systems` file, 590
- telnet command, secure NFS and, 232
- template files (PPP)
 - `/etc/ppp/myisp-chat.tmpl`, 538
 - `/etc/ppp/options.tmpl`, 528
 - `/etc/ppp/peers/myisp.tmpl`, 534
 - `options.ttya.tmpl`, 530
- template files for PPP, list of templates, 461
- temporary (TM) UUCP data files, 618
- testing
 - mail aliases, 345
 - mail configuration, 344
 - mail connections to other systems, 347
 - rule sets, 346
- testing packet reliability, 681
- Th Time field entry, 588
- threshold values, 91
- tilde (~)
 - abbreviated path names, 673, 674
 - `rcp` command syntax, 676, 677
- time
 - synchronizing with another system, 57
- Time field of `Systems` file, 588, 607
- time synchronization, 231
- Timeout option
 - `sendmail` command, 399, 404
- timeouts (SLP), 277, 284
- TLI dialer type, 596
- TLI network, 596
- TLIS dialer type, 596
- TM UUCP temporary data files, 618
- tokens (dialer token pairs), 596, 598
- TOO MANY LOCKS message, 621
- TOO MANY SAVED C FILES message, 620
- transfer speed for UUCP communication
 - link, 589, 590, 595
- Transport Level Interface Network (TLI), 596
- transport protocol negotiation, 222
- transport setup problem, error message, 192
- troubleshooting
 - autofs, 189
 - avoiding mount point conflicts, 172
 - autofs (Continued)
 - error messages generated by `automount`
 - `-v`, 189
 - miscellaneous error messages, 190
 - mail aliases, 345
 - mail connections to other systems, 347
 - mail services, 344
 - MAILER-DAEMON messages and, 348
 - networks, 686, 689
 - NFS
 - determining where NFS service has failed, 186
 - hung programs, 194
 - miscellaneous error messages, 192
 - remote mounting problems, 183, 193
 - server problems, 183
 - strategies, 182
 - rule sets, 346
 - undelivered mail, 345
 - UUCP, 583, 622
 - ASSERT error messages, 585, 619, 621
 - checking basic information, 585
 - checking error messages, 585, 622
 - checking `Systems` file, 584
 - commands for troubleshooting, 585
 - debugging transmissions, 583, 584
 - faulty modem or ACU, 583
 - STATUS error messages, 585, 621, 622
- troubleshooting PPP
 - common problems, 506
 - authentication, 519
 - chat scripts, 515, 516
 - for networks, 510
 - general communications, 512
 - leased-line links, 518
 - serial lines, 517
 - with the PPP configuration, 513
 - obtaining diagnostics, 507, 508
 - task map, 505
- truss command, 222
- trusted callers, 435
 - configuring for CHAP authentication, 493
 - configuring for PAP authentication, 484, 487, 488, 489
- trusted-users file, 367, 417
- TrustedUser option, `sendmail` command, 399

- trusting network environment
 - remote login
 - authentication process, 659
 - process after logging in, 662
- Tu Time field entry, 588
- tuning SLP performance, 272
- tunnel
 - definition (PPP), 437
 - example configuration, 456, 457
 - participants, 437
 - task maps for configuring, 497
- turning off
 - CLOCAL flag, 592
 - echo checking, 591, 601
- turning on
 - CLOCAL flag, 592
 - echo checking, 591, 601
 - enabling dialback through chat script, 592
- Type field
 - Devices file, 594, 595
 - Systems file, 589

U

- U option, sendmail command, 392
- UA, requests, 273
- UAs (SLP), 260, 284, 286
- UDP, NFS version 3 and, 147
- UDP protocol, 684
- UDP/TCP unicast (SLP), 287
- UMASK parameter, /etc/default/nfslogd file, 199
- umount command
 - autofs and, 145
 - description, 208
- umountall command, 209
- uname -n command, 605
- undelivered messages, troubleshooting, 345
- underscore (_) in mailbox names, 360
- unicast routing (SLP), 287
 - disabled, 289
- UNIX authentication, 229, 230
- unmounting
 - autofs and, 145, 242
 - examples, 209, 210
 - groups of file systems, 209

- UnsafeGroupWrites option, sendmail command, 400
- unshare command, 214
- unshareall command, 215
- unsharing file systems
 - unshare command, 214
 - unshareall command, 215
- updetach option (PPP), 534
- URL service types, WebNFS and, 167
- UseErrorsTo option, sendmail command, 400
- UseMSP option, sendmail command, 399
- Usenet, 571, 587
- user agent (SLP), 268
- User-job-grade field of Grades file, 614
- User keyword of Permit-type field, 616
- user names
 - current user, 674
 - direct or indirect logins (rlogin), 661
 - finding users logged in to remote system, 664
- user names, mailbox names and, 360
- user option (PPP), 486, 533
- /usr/bin/aliasadm command, 366
- /usr/bin/cu command
 - checking modems or ACUs, 583
 - description, 574
 - multiple or different configuration files, 575, 604
 - printing Systems lists, 605
- /usr/bin directory, contents of, 366
- /usr/bin/mail command, 366
- /usr/bin/mailcompat filter, 366
- /usr/bin/mailq command, 366
- /usr/bin/mailstats command, 366
- /usr/bin/mailx command, 366
- /usr/bin/mconnect command, 366
- /usr/bin/mconnect command, 347
- /usr/bin/ncab2clf command, 51
- /usr/bin/praliases command, 366
- /usr/bin/rmail command, 366
- /usr/bin/uucp command
 - debugging transmissions, 584
 - description, 574
 - home directory of login ID, 573
 - permissions for forwarding operation, 613
 - uucico execution by, 572
- /usr/bin/uulog command, 573, 585

- /usr/bin/uupick command, 574, 583
- /usr/bin/uustat command, 574, 583
- /usr/bin/uuto command
 - description, 574
 - removing public directory files, 583
 - uucico execution by, 572
- /usr/bin/uux command
 - description, 574
 - uucico execution by, 572
- /usr/bin/vacation command, 366, 374
- /usr directory, mounting by diskless clients, 145
- /usr/dt/bin/dtmail mail user agent, 370
- /usr/kvm directory, mounting by diskless clients, 145
- /usr/lib directory, contents of, 367
- /usr/lib/inet/xntpd daemon, 58
- /usr/lib/mail/cf directory, 368
- /usr/lib/mail/cf/main.mc file, 368, 417
- /usr/lib/mail/cf/main-v7sun.mc file, 417
- /usr/lib/mail/cf/makefile file, 368
- /usr/lib/mail/cf/submit.mc file, 368
- /usr/lib/mail/cf/subsidiary.mc file, 368, 418
- /usr/lib/mail/cf/subsidiary-v7sun.mc file, 418
- /usr/lib/mail directory, contents of, 368
- /usr/lib/mail/domain directory, 368
- /usr/lib/mail/domain/generic.m4 file, 368
- /usr/lib/mail/domain/solaris-antispam.m4 file, 369
- /usr/lib/mail/domain/solaris-generic.m4 file, 369
- /usr/lib/mail/feature directory, 369
- /usr/lib/mail.local mailer, 367
- /usr/lib/mail/m4 directory, 369
- /usr/lib/mail/mailer directory, 369
- /usr/lib/mail/ostype directory, 369
- /usr/lib/mail/ostype/solaris2.m4 file, 369
- /usr/lib/mail/ostype/solaris8.m4 file, 369
- /usr/lib/mail/ostype/solaris2.ml.m4 file, 369
- /usr/lib/mail/ostype/solaris2.pre5.m4 file, 369
- /usr/lib/mail/README file, 368
- /usr/lib/mail/sh/check-hostname script, 369
- /usr/lib/mail/sh/check-permissions script, 369
- /usr/lib/mail/sh directory, 369
- /usr/lib/nca_addr.so library, 51
- /usr/lib/net/ncaconfd command, 51
- /usr/lib/nfs/nfslogd daemon, enabling NFS server logging, 156
- /usr/lib/sendmail command, 367
- /usr/lib/smrsh command, 368
- /usr/lib/uucp/uuccheck command, 573, 585
- /usr/lib/uucp/uucleanup command, 573
- /usr/lib/uucp/Uutry command, 573, 583, 584
- /usr/ntp/ntpstats directory, 58
- /usr/sbin/editmap command, 370
- /usr/sbin/etrn script, 370
- /usr/sbin/in.comsat daemon, 370
- /usr/sbin/inetd daemon, in.uucpd invoked by, 573
- /usr/sbin/makemap command, 370
- /usr/sbin/mount command
 - See mount command
- /usr/sbin/newaliases link, 370
- /usr/sbin/ntpdate command, 58
- /usr/sbin/ntpq command, 58
- /usr/sbin/ntptrace command, 58
- /usr/sbin/shareall command
 - See shareall command
 - automatic file-system sharing, 153
 - enabling WebNFS service, 155
- /usr/sbin/showmount command, 216
- /usr/sbin/sppptun command
 - definition, 556
 - example, 556, 557
- /usr/sbin/syslogd command, 370
- /usr/sbin/unshareall command, 215
- /usr/sbin/xntpd command, 58
- uuccheck command, 573, 585
- uucico daemon
 - adding UUCP logins, 578
 - description, 572
 - Dialcodes file and, 604
 - maximum simultaneous executions, 575, 617

- uucico daemon (Continued)
 - multiple or different configuration files, 575, 587, 604
 - printing Systems lists, 605
 - Systems file and, 587
 - uusched daemon and, 573
 - Uutry command and, 573
- uucleanup command, 573
- UUCP
 - "login shell", 572
 - administrative commands, 573
 - administrative files, 618, 619
 - callback option, 608, 609
 - configuring
 - adding UUCP logins, 578
 - running UUCP over TCP/IP, 581
 - daemons
 - overview, 572, 573
 - database files, 574, 617
 - asppp configuration, 575, 576
 - basic configuration files, 575, 576
 - description, 574, 575
 - multiple or different files, 575, 587, 604
 - description, 571, 587
 - directories
 - administration, 573
 - error messages, 585
 - public directory maintenance, 583
 - displaying log files, 573
 - file transfers
 - daemon, 572
 - permissions, 606, 608
 - troubleshooting, 583, 584
 - work files C., 618, 619
 - forwarding operation, 613
 - hardware configurations, 571
 - log files
 - cleanup, 580
 - displaying, 573
 - logins
 - adding, 578
 - privileges, 611
 - mail accumulation, 582
 - maintenance, 582, 583
 - node name
 - alias, 575, 607
 - remote computer, 588, 605
 - overriding parameters manually, 613
- UUCP (Continued)
 - passive mode, 607
 - polling remote computers, 575, 613
 - privileged logins and passwords, 611
 - public directory maintenance, 583
 - remote execution
 - commands, 606, 609, 612
 - daemon, 573
 - work files C., 618, 619
 - security
 - COMMANDS option of Permissions file, 609, 610
 - setting up, 582
 - sticky bit for public directory files, 583
 - VALIDATE option of Permissions file, 610, 612
 - shell scripts, 579, 581
 - Solaris version, 571, 587
 - spool
 - clean-up command, 573
 - job grade definitions, 614, 616
 - scheduling daemon, 573
 - STREAMS configuration, 616
 - transfer speed, 589, 590, 595
 - troubleshooting, 583, 622
 - ACU faulty, 583
 - ASSERT error messages, 585, 619, 621
 - checking basic information, 585
 - checking error messages, 585, 622
 - checking Systems file, 584
 - commands for troubleshooting, 585
 - debugging transmissions, 583, 584
 - modem faulty, 583
 - STATUS error messages, 585, 621, 622
 - user commands, 574
- UUCP (UNIX-to-UNIX Copy command)
 - mailers, 357
 - testing the connection, 345
- uucp command
 - debugging transmissions, 584
 - description, 574
 - home directory of login ID, 573
 - permissions for forwarding operation, 613
 - uucico execution by, 572
- uucppublic directory maintenance, 583
- uudemon.admin shell script, 580
- uudemon.cleanup shell script, 580
- uudemon.crontab file, 579

- uudemon.hour shell script
 - description, 580
 - uusched daemon execution by, 573
 - uuxqt daemon execution by, 573
 - uudemon.poll shell script, 580, 613
 - uudirect keyword of DTP field, 596
 - uulog command, 573, 585
 - uname command, 585
 - uupick command
 - description, 574
 - removing public directory files, 583
 - uusched daemon
 - description, 573
 - maximum simultaneous executions, 575, 617
 - uudemon.hour shell script call, 580
 - uustat command
 - checking modems or ACUs, 583
 - description, 574
 - uudemon.admin shell script for, 580
 - uuto command
 - description, 574
 - removing public directory files, 583
 - uucico execution by, 572
 - Uutry command, 573, 583
 - Uutry command, 584
 - uux command
 - description, 574
 - uucico execution by, 572
 - uuxqt daemon
 - description, 573
 - maximum simultaneous executions, 575, 617
 - uudemon.hour shell script call, 580
- V**
- V option, umount command, 208
 - v option
 - automount command, 189
 - uucheck command, 585
 - vacation command, 366, 374
 - VALIDATE option of Permissions file, 610, 612
 - COMMANDS option, 609, 610
 - /var/adm/exacct directory, 84
 - /var/mail directory, 312, 313
 - automatic mounting of, 317
 - mail client configuration and, 317
 - /var/mail file, 360
 - /var/nca/log file, 51
 - /var/run/nca_httpd_1.door file, 51
 - /var/run/sendmail.pid file, 370
 - /var/spool/clientmqueue directory, 370
 - /var/spool/mqueue directory, 370
 - /var/spool/uucppublic directory
 - maintenance, 583
 - /var/uucp/.Admin/errors directory, 585
 - /var/uucp/.Status directory, 585
 - variables in map entries, 244, 245
 - vendor setting, specifying in sendmail.cf file, 354
 - verifiers
 - description, 230
 - UNIX authentication, 230
 - verifying, remote system operation, 663
 - version 2 NFS protocol, 146
 - version 3 NFS protocol, 146
 - version level, specifying in sendmail.cf file, 354
 - version negotiation, 222
 - vfstab file
 - automount command and, 238
 - description, 198
 - enabling client side failover, 160
 - mounting by diskless clients, 145
 - mounting file systems at boot time, 157
 - NFS servers and, 158
 - nolargefiles option, 160
 - virtual hosts, setting up, 322
 - VIRTUSER_DOMAIN_FILE() m4 configuration macro, 409
 - VIRTUSER_DOMAIN() m4 configuration macro, 409
 - virtuser_entire_domain FEATURE() declaration, 411
- W**
- warm start, rpcbind daemon, 187
 - WARNING: mountpoint already mounted on message, 190
 - We Time field entry, 588
 - WebNFS service
 - browsing, 167
 - description, 227

Composed September 5, 2002

- WebNFS service (Continued)
 - enabling, 154
 - firewalls and, 167
 - overview, 148
 - planning for, 166
 - security negotiations and, 149
 - task map, 165
 - URL service types and, 167
- weighting of servers in maps, 244
- wide area network (WAN)
 - Usenet, 571, 587
- Wk Time field entry, 588
- work (C.) UUCP files
 - cleanup, 580
 - description, 618, 619
- working directory, definition for rcp
 - command, 674
- write errors, NFS and, 146
- WRITE option of Permissions file, 608
 - NOWRITE option, 608
- WRONG MACHINE NAME message, 622
- WRONG ROLE message, 620
- WRONG TIME TO CALL message, 621

X

- X. UUCP execute files
 - cleanup, 580
 - description, 619
 - uuxqt execution, 573
- XMV ERROR message, 621
- xntpd daemon, 58
- xntpd daemon
 - starting, 56
- xntpd script, 58
- xntpd command, 58
- xonxoff option (PPP), 472
- XscriptFileBufferSize option, sendmail
 - command, 399

Composed September 5, 2002