

IRIX® Admin:
Networking and Mail

007-2860-012

CONTRIBUTORS

Written by Arthur Evans and Jeffrey B. Zurschmeide

Updated by Pam Sogard, Helen Vanderberg, Bob Bernard, Terry Schultz, Julie Boney,
and Ken Jones

Production by Karen Jacobson and Ken Jones

Engineering contributions by Michel Bourget, Scott Henry, Carlin Otto, Kam Kashani, Chris
Wagner, Paul Mielke, Robert Stephens, Joe Yetter, Gretchen Helms, John Schimmel, Robert
Mende, Vernon Schryver, Michael Nelson, and Landon Noll

Illustrations by Dany Galgani

COPYRIGHT

© 1996 —2005 Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in
third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or
create derivative works from the contents of this electronic documentation in any manner, in
whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with
restricted rights (except as to included open/free source) as specified in the FAR 52.227-19
and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation
of worldwide intellectual property laws, treaties and conventions. This document is provided
with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, Challenge, IRIS, IRIX, Onyx, and Origin are registered
trademarks and 4DDN, 4DLT, FDDI Visualyzer, IRIX NetWorker, IRIS 4D, and NetVisualyzer
are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

DSI is a trademark of Digicom Systems, Inc. FLEXIm is a trademark of GLOBEtrrotter Software,
Inc. Hayes is a registered trademark of Hayes Microcomputer Products, Inc. IBM 3270 is a
trademark of International Business Machines, Inc. Intel is a registered trademark of Intel
Corporation. Macintosh is a registered trademark of Apple Computer Corporation. MS-DOS is
a registered trademark of Microsoft Corporation. Sun and RPC are registered trademarks of Sun
Microsystems, Inc. Tektronix is a trademark of Tektronix, Inc. Telebit is a registered trademark
of Telebit Corporation. Robotics is a registered trademark of U. S. Robotics, Inc. UNIX and X
Window System are registered trademarks of The Open Group in the United States and other
countries. ZyXEL is a trademark of ZyXEL.

What's New in This Guide

This rewrite of the *IRIX Admin Networking and Mail* guide includes the following changes:

- Updated section “Setting Up an Anonymous FTP Account” in Chapter 3.
- Miscellaneous corrections

Record of Revision

Version	Description
006	February 2000 Incorporates information for the IRIX 6.5.7 release.
007	February 2002 Incorporates information for the IRIX 6.5.15 release.
008	February 2003 Incorporates information for the IRIX 6.5.19 release.
009	Novembor 2003 Incorporates information for the IRIX 6.5.22 release.
010	February 2004 Incorporates information for the IRIX 6.5.23 release.
011	May 2004 Incorporates information for the IRIX 6.5.24 release.
012	October 2005 Miscellaneous corrections.

Internet Protocol Version 6 18
IPv6 Overview 19
IPv6 Addressing 20
Format of Internet Protocol IP Version 6 Addresses 21
IPv6 Header Format 22
Installing and Configuring IPv6 24
Domain Names 24
Obtaining a Domain Name 25
About Subdomains 25
About Internet Connections 25
Before Connecting to the Internet 26
About Local Network Information Centers. 27
Online Information Sources 30
About Name-to-Address Mapping 32
/etc/hosts Database 33
Domain Name System 33
Network Information Service (NIS) 33
Guidelines for a Subnetwork 34
IP Address Allocation 36
About Network Security 37
Common Network Applications 37
About Electronic Mail. 37
About Network File System (NFS) 37
3. Setting Up a Network 39
Configuring a System for a Network. 39
Attaching Your Station to an Ethernet Network 40
Checking Your Ethernet Connection 40
Checking the Network Software Configuration 43
About the Hosts Database 43
Modifying the Hosts Database 45
Naming Your Station 46
Testing Your Network Connectivity. 46

Setting Up a Router	47
Configuring a Router With Two Interfaces	47
Configuring a Router With More Than Two Interfaces	48
Configuring Routing Behavior	49
Turning On Multicast Routing	50
About Multicast Packets Forwarding	51
Setting Up Tunnels to Support Multicast Packets	53
Updating /etc/rpc for NIS Users	54
Subnetting a Network.	55
Setting the Netmask	55
Rebooting the Station.	56
Modifying the Network Interface Configuration With the /etc/config/netif.options File	56
Modifying the Interface Name in the /etc/config/netif.options File	58
Modifying the Interface Address in the /etc/config/netif.options File	59
Assigning IP Aliases	60
Changing Network Parameters in the ifconfig-#.options File	62
Configuring /etc/gateways Files for Networks That Do Not Support Broadcast or Multicast	64
/etc/gateways File Format	64
Examples of /etc/gateways Files	66
Configuring Multiple Network Interfaces	67
Dynamic Host Configuration With Proclaim	68
DHCP Server Basic Operation	69
Configuring the DHCP Server	69
Configuring the DHCP Relay Agent	83
Setting up the DHCP Proclaim Client	85
Limitations of DHCP	90
Creating a Local Network Script.	91
Turning On Remote Access Logging	91
Setting Up Network-Wide Services	92
Setting Up an Anonymous FTP Account	92
Setting Up a Password-Protected FTP Account	97
About InfoSearch File Servers	99
Setting up Time Synchronization Service	103

Accessing Network Services through Internet Gateway	104
Reserving Resources With RSVP	105
Installing RSVP	106
Troubleshooting RSVP	107
Troubleshooting Your Ethernet Connection.	108
Troubleshooting Cable Problems	108
Troubleshooting Late-Collision Problems	109
Troubleshooting Packet Size Problems	110
Troubleshooting Server Contact Problems	110
Checking Additional Network Interfaces	111
Configuring a Dual-port Gigabit Ethernet Board	112
How To Change Your Primary Network Interface	112
4. Introducing Network Management	115
Resources for Network Management	115
About Network Management Functions	116
About Network Startup and Shutdown	117
Network Initialization Process	117
Network Shutdown Process	119
Network Management Tools	119
Interpreting Network Statistics	123
Testing Network Connectivity With ping	123
Measuring Network Throughput With tcp	124
Collecting Network Statistics With netstat	126
Network Tuning Information	127
About Setting MTU Sizes	128
About Setting Packet Forwarding	128
About Setting Window Sizes	128
HTTP Considerations	129

Troubleshooting Poor Network Performance129
Troubleshooting Hardware Problems to Improve Network Performance129
Troubleshooting Network Configuration to Improve Network Performance130
Troubleshooting Network Daemons to Improve Network Performance131
Decreasing Packet Size to Improve Network Performance131
Kernel Configuration for Better Network Performance132
5. SLIP and PPP133
About SLIP and PPP134
Setting Up a SLIP or PPP Connection: General Procedures135
Verifying the SLIP and PPP Software136
Installing the SLIP and PPP Software137
Selecting a Modem137
IP Addresses for SLIP and PPP Clients137
Configuring a System for Dial-Out138
Configuration Files for Dial-Out138
Sample SLIP Configuration for Dial-Out142
Sample PPP Configuration for Dial-Out142
Configuring a System for Dial-In144
Configuring /etc/passwd for Dial-In With SLIP144
Configuring /usr/etc/remoteslip for Dial-In With SLIP.145
Configuring /etc/passwd for Dial-In With PPP146
SLIP and PPP Routing and Address Allocation146
About Proxy-ARP Routing for SLIP Connections.147
Setting Up SLIP/PPP Subnets for Client Addresses149
About Connected SLIP or PPP Networks149
Using Dynamic Address Allocation With PPP150
Configuring a Bidirectional Link.150
Starting SLIP or PPP at Boot Time150
About Demand Dialing151
Setting Up Demand Dialing151
NFS Over SLIP or PPP152
File Transfer Over SLIP or PPP152
Troubleshooting SLIP and PPP Links152

6. BIND Name Server	155
About Domain Name Service	156
BIND Servers and Clients	158
BIND Master Servers	159
BIND Slave and Forwarding Servers	159
BIND Caching-Only Server	160
BIND Clients	160
BIND Configuration Files	161
BIND Boot File	162
BIND named.hosts File	164
BIND named.rev File	165
BIND localhost.rev File	165
BIND root.cache File	165
BIND /etc/config/named.options File	165
hoResolution With /etc/resolv.conf.	166
Setting Up a BIND Configuration	168
Configuring the Primary BIND Server	169
Configuring the Secondary BIND Server	173
Configuring a Caching-Only BIND Server	174
Configuring the Forwarding BIND Server	175
Configuring a Slave BIND Server	176
Configuring the BIND Client.	177
Managing the BIND Environment	177
Adding a New BIND Station	178
Deleting a BIND Station	178
Adding Another BIND Domain	178
named Reload Script	178
named Restart Script	179
Debugging named	179
SYSLOG Error Messages	180
Debugging Name Servers With the nslookup Command.	181

Introduction to BIND Version 9181
BIND Version 9 Overview182
Migrating to BIND 9183
Implementation differences for BIND 9 on IRIX184
7. Unified Name Service185
About Unified Name Service.186
Overview of UNS Operations186
How UNS Works With NIS189
About UNS and the NIS Database192
How UNS Works With BIND192
How UNS Works With NFS193
How UNS Works With LDAP194
Namespace Format195
UNS Configuration File196
Setting Up a UNS Configuration.197
UNS Protocol Libraries197
Files Library199
NIS Protocol Library199
DNS Protocol Library199
MDBM Protocol Library199
NDBM Protocol Library200
Berkeley DB Protocol Library200
Nisserv Protocol Library200
LDAP Protocol Library201
UNS File Structure Attributes201
Setting Domain Attributes201
Setting Table Attributes202
Setting Library Attributes202
Setting File Attributes202
Querying Attributes202
Cache Tuning203

Troubleshooting nsd	203
General Approach to Troubleshooting	203
Deciphering nsd Signals	204
Verifying /ns/.local	204
8. UUCP	205
Choosing TCP/IP or UUCP	206
Hardware Requirements for UUCP	207
UUCP Commands	207
UUCP User Program Commands	207
UUCP Administrative Program Commands	208
UUCP Daemons	209
Supporting UUCP Databases	210
UUCP Devices File	211
UUCP Dialers File	216
UUCP Systems File	219
UUCP Dialcodes File	223
UUCP Permissions File	223
UUCP Poll File	231
UUCP Sysfiles File	231
Other UUCP Files	232
UUCP Administrative Files	233
Setting Up UUCP	235
Determining the Remote and Local Stations for a UUCP Connection	235
Making the Physical Connection for UUCP	236
Configuring the Local Station for UUCP	236
Configuring the Remote Station for UUCP	240
Testing the UUCP Connection	243
Setting Up UUCP on a TCP/IP Connection	245
UUCP Error Messages	247
ASSERT Error Messages	247
STATUS Error Messages	249

9.	IRIX sendmail253
	About the Mail System254
	Overview of sendmail.255
	System Organization256
	Structure of sendmail257
	About the sendmail Components259
	sendmail Daemon.259
	sendmail Accounts260
	sendmail Scripts260
	sendmail Related Files and Directories.262
	Aliases Database for sendmail265
	Building the sendmail Aliases Database266
	Testing the sendmail Aliases Database.268
	sendmail Alias Database Problems268
	sendmail List Owners269
	Configuring sendmail.269
	sendmail.cf File269
	sendmail.mc and sendmail.cf Files270
	Managing sendmail275
	Starting the sendmail Daemon276
	Listing the sendmail Message Queue276
	Forcing the sendmail Message Queue276
	Redirecting Mail With the .forward File278
	About sendmail MX Records.279
10.	SGI Infinite Network Bandwidth281
	SGI Infinite Network Bandwidth Overview282
	New Terminology Used in This Guide.282
	Definition of SGI Infinite Network Bandwidth283
	Hardware Requirements.283
	Configuring a Network Stripe286
	Verifying a Network Stripe289

Stripe Interface MTUs	290
Default MTU Calculation	291
MTUs and Performance	291
Network Stripe Monitoring	293
Tuning a Network Stripe Interface	294
Stripe Interface System Tunable Parameters	295
<i>ifconfig</i> Parameters	295
Binding and Placement	297
<i>tg</i> Gigabit Ethernet Interfaces	297
The Reorder Thread	298
Applications	298
Application I/O Sizes	299
Troubleshooting SGI Infinite Network Bandwidth Software	300
Testing Connectivity	300
Diagnosing Performance Problems	302
BIND Standard Resource Record Syntax	306
About TTLs in BIND Resource Records	306
About Special Characters in BIND Resource Records	307
Specifying \$INCLUDE in BIND Resource Records	307
Specifying \$ORIGIN in BIND Resource Records	308
Specifying SOA—Start of Authority in BIND Resource Records	308
Specifying NS—Name Server in BIND Resource Records	309
Specifying A—Address in BIND Resource Records	310
Specifying HINFO—Host Information in BIND Resource Records	310
Specifying WKS—Well-Known Services in BIND Resource Records	310
Specifying CNAME—Canonical Name in BIND Resource Records	311
Specifying PTR—Domain Name Pointer in BIND Resource Records	311
Specifying MB—Mailbox in BIND Resource Records	311
Specifying MR—Mail Rename Name in BIND Resource Records	312
Specifying MINFO—Mail Information in BIND Resource Records	312
Specifying MG—Mail Group Member in BIND Resource Records	312
Specifying MX—Mail Exchanger in BIND Resource Records	312
Specifying RP—Responsible Person in BIND Resource Records	313

Specifying TXT—Text in BIND Resource Records314
Specifying the sendmail Debugging Flags315
sendmail Debugging Traceflags316
Index319

List of Figures

Figure 1-1	Ethernet Network Attachment	2
Figure 1-2	Serial Line Network	3
Figure 2-1	Heterogeneous Network With Wide-Area Connections	12
Figure 2-2	Format of Internet Protocol (IP) Version 4 Addresses	16
Figure 2-3	Format of Internet Protocol (IP) Version 6 Addresses	21
Figure 2-4	IPv6 Prefix Bits and Interface Identifier Bits.	22
Figure 2-5	IPv6 Header Format and Data Portion of a Datagram	23
Figure 2-6	Subnetted Class B Address	35
Figure 3-1	Network With Multicast Routers	52
Figure 3-2	A Tunnel Between Networks A and C	53
Figure 3-3	ProclaimServerMgr Graphical User Interface	71
Figure 3-4	ProclaimServerMgr File Button Pulldown Menu	71
Figure 3-5	ProclaimRelayMgr Graphical Interface (GUI)	84
Figure 6-1	Partial View of Domain Name Space.157
Figure 6-2	Example BIND Configuration.169
Figure 7-1	Action of the nsd Daemon With Name Service Protocols187
Figure 7-2	Partial View of Dynamic UNS Files188
Figure 7-3	Historical Operation of NIS Name Lookups190
Figure 7-4	Operation of the nsd Daemon With NIS.191
Figure 7-5	Action of ns_lookup in Selection of Protocol Library198
Figure 9-1	Layers of TCP/IP Mail Software255
Figure 9-2	sendmail System Structure258

List of Tables

Table 1-1	Standard Networking Software	5
Table 1-2	Optional Networking Products	5
Table 2-1	Network Device Characteristics	9
Table 2-2	Network Information Centers.	27
Table 3-1	Variables for the netif.options File	57
Table 3-2	IRIX Time Synchronization Services	103
Table 3-3	IRIX Time Synchronization Services Default Configuration Parameters. 103	
Table 3-4	Supported Cards for RSVP	106
Table 6-1	BIND Server Configurations	158
Table 6-2	named Database Files	162
Table 7-1	Protocols With Historically Supported Services	189
Table 7-2	UNS Files and Their Purposes	195
Table 7-3	Cache Tuning Parameters	203
Table 8-1	Comparison of TCP/IP and UUCP	206
Table 8-2	UUCP Escape Sequences	217
Table 8-3	Three-Wire Null-Modem Pinning Configuration	236
Table 8-4	Assert Error Messages	247
Table 8-5	STATUS Error Messages	249
Table 10-1	Throughput and Estimated Socket Buffer Size	296

About This Guide

This guide explains how to set up and maintain a network of SGI workstations and servers. It includes information on TCP/IP networking, including SLIP and PPP, UUCP networking, and configuring the sendmail mail transfer agent.

The standard network communications software that runs on SGI workstations is derived from the networking software in the 4.3BSD UNIX releases from the University of California at Berkeley and the Sun Microsystems RPC (remote procedure call) system. The IRIX operating system implements the Internet Protocol suite and UNIX domain sockets using the 4.3BSD UNIX socket mechanism. The system also supports access to the underlying network media by means of raw sockets.

What This Guide Contains

IRIX Admin: Networking and Mail contains the following chapters:

- Chapter 1, “About Networking Products,” discusses SGI standard hardware and software networking products and describes the standard software configuration (files, daemon, processes).
- Chapter 2, “Planning a Network,” provides insight into planning a network. It includes internet addressing, the *hosts* database file, when to use certain applications, how to subnet a network, security issues, and heterogeneous network considerations.
- Chapter 3, “Setting Up a Network,” describes, through example, the process of configuring a network (homogeneous and heterogeneous), how to set up a router, and basic troubleshooting advice.
- Chapter 4, “Introducing Network Management,” describes the various tools available for managing a network, including backup strategies, performance issues, and fault isolation.
- Chapter 5, “SLIP and PPP,” describes the features and functions of SLIP and details how to connect two stations using SLIP.

- Chapter 6, “BIND Name Server,” provides an overview of the Berkeley Internet Name Domain (BIND) server, also known as `named`. It also provides an example setup procedure and general information on managing and troubleshooting BIND.
- Chapter 7, “Unified Name Service,” provides an overview of the Unified Name Server, `nsd`. It considers the interconnections with other name services and contains general information on troubleshooting UNS.
- Chapter 8, “UUCP,” compares TCP/IP and UUCP and describes the features and functions of the UUCP networking utilities. It also provides a setup example and information about common UUCP error messages.
- Chapter 9, “IRIX sendmail,” provides an overview of the mail system, the `sendmail` program, and the `alias` database. It contains a planning checklist and a setup example for various `sendmail` configurations.
- Chapter 10, “SGI Infinite Network Bandwidth,” describes network striping that provides point-to-point bandwidth aggregation over multiple interfaces. Network striping provides higher throughput transparently to network-based applications such as file transfer protocol (FTP), Remote Procedure Call (RPC), network file system (NFS) and user socket-based applications. It is described in detail in a new Chapter 10.
- Appendix A, “BIND Standard Resource Record Format”, provides detailed information about all standard resource record formats used in BIND configuration files.
- Appendix B, “IRIX sendmail Reference”, provides a concise reference to `sendmail` as it is implemented from the `sendmail` standard.

Conventions Used in This Guide

These type conventions and symbols are used in this guide:

<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, pathnames, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[]	Brackets enclose optional portions of a command or directive line.

manpage(x)	Man page section identifiers appear in parentheses after man page names.
""	(Double quotation marks) References in text to document section titles
#	IRIX shell prompt for the superuser (root)
%	IRIX shell prompt for users other than superuser

Additional Resources

The *IRIX Admin* guides are available through the InfoSearch online viewing system. They are also available on the World Wide Web at <http://docs.sgi.com>. The set comprises these volumes:

- *IRIX Admin: Software Installation and Licensing*—Explains how to install and license software that runs under IRIX, the Silicon Graphics implementation of the UNIX operating system. Contains instructions for performing miniroot and live installations using Inst, the command line interface to the IRIX installation utility. Identifies the licensing products that control access to restricted applications running under IRIX and refers readers to licensing product documentation.
- *IRIX Admin: System Configuration and Operation*—Lists good general system administration practices and describes system administration tasks, including configuring the operating system; managing user accounts, user processes, and disk resources; interacting with the system while in the PROM monitor; and tuning system performance.
- *IRIX Admin: Disks and Filesystems*—Explains disk, filesystem, and logical volume concepts. Provides system administration procedures for SCSI disks, XFS and EFS filesystems, XLV logical volumes, and guaranteed-rate I/O.
- *IRIX Admin: Networking and Mail*—Describes how to plan, set up, use, and maintain the networking and mail systems, including discussions of sendmail, UUCP, SLIP, and PPP.
- *IRIX Admin: Backup, Security, and Accounting*—Describes how to back up and restore files, how to protect your system's and network's security, and how to track system usage on a per-user basis.
- *IRIX Admin: Resource Administration*—Provides an introduction to system resource administration and describes how to use and administer various IRIX resource management features, such as IRIX process limits, IRIX job limits, Array Services, Comprehensive System Accounting (CSA), Cpuset System, and Miser.

- *IRIX Admin: Peripheral Devices*—Describes how to set up and maintain the software for peripheral devices such as terminals, modems, printers, and CD-ROM and tape drives.
- *IRIX Admin: Selected Reference Pages* (not available in InfoSearch)—Provides concise reference page (manual page) information on the use of commands that may be needed while the system is down. Generally, each reference page covers one command, although some reference pages cover several closely related commands. Reference pages are available online through the `man` command.
- *Performance Co-Pilot for IRIX Advanced User's and Administrator's Guide* —Describes the Performance Co-Pilot (PCP) suite of advanced performance management applications for SGI workstations and servers. PCP provides a systems-level set of tools that cooperate to deliver distributed, integrated performance monitoring and performance management services spanning the hardware platform, operating system, multiple layered services, and end-user applications.

Internet Request For Comment documents are available from the Internet Network Information Center (InterNIC) at the following address:

Network Solutions
Attn: InterNIC Registration Services
505 Huntmar Park Drive
Herndon, VA 22070
Phone: 1-800-444-4345 or 1-703-742-4777

Internet Request For Comment documents are also available by anonymous `ftp` from various sites, such as `ftp.ds.internic.net`.

RFC 2461, *Neighbor Discovery for IP Version 6 (IPv6)*

RFC 2373, *IPv6 Addressing Architecture*

Abitz, P, Liu, C., *DNS and BIND* (Sebastopol, CA: O'Reilly & Associates, Inc.).

Braden, R. "Requirements for Internet Hosts." *Internet Request For Comment 1112* (1989).

Comer, D. E., *Internetworking with TCP/IP Volume 1*. (Englewood Cliffs, NJ: Prentice-Hall, 1995).

Costales, B. with Allman, E, *sendmail*. (Sebastopol, CA: O'Reilly & Associates, Inc., 1997).

Deering, S. "Host Extensions for IP Multicasting." *Internet Request For Comment 1112* (1989).

Everhart, C., Mamakos, L., Ullmann, R., Mockapetris, P. "New DNS RR Definitions." *Internet Request For Comment 1183* (1990).

Held, G., *LAN Management with SNMP and RMON*. (J. Wiley and Sons, 1996).

Huitema, C., *Routing in the Internet*. (Englewood Cliffs, NJ: Prentice-Hall, 1995).

Hunt, C., *TCP/IP Network Administration*. (Sebastopol, CA: O'Reilly & Associates, Inc., 1992).

Leinwand, A., Conroy, K.F., *Network Management - A Practical Perspective*. (Addison Wesley, 1996).

Lottor, M. "Domain Administrator's Guide." *Internet Request For Comment 1033* (1987).

Lottor, M. "TCP Port Service Multiplexer (TCPMUX)." *Internet Request For Comment 1078* (1988).

Loukides, M., *System Performance Tuning*. (Sebastopol, CA: O'Reilly & Associates, Inc., 1990).

Mockapetris, P. "DNS Encoding of Network Names and Other Types." *Internet Request For Comment 1101* (1989).

Mockapetris, P. "Domain Names – Concept and Facilities." *Internet Request For Comment 1034* (1987).

Mockapetris, P. "Domain Names – Implementation and Specification." *Internet Request For Comment 1035* (1987).

Mogul, J., Postel, J. "Internet Standard Subnetting Procedure." *Internet Request for Comment 950* (1985).

Partridge, C. "Mail Routing and The Domain System." *Internet Request For Comment 974* (1986).

Stahl, M. "Domain Administrator's Guide." *Internet Request For Comment 1032* (1987).

Stern, H., *Managing NFS and NIS*. (Sebastopol, CA: O'Reilly & Associates, Inc., 1991).

Stevens, W. R., *TCP/IP Illustrated, Volume 1*. (Addison Wesley, 1996).

Obtaining Publications

To obtain SGI documentation, go to the SGI Technical Publications Library at <http://docs.sgi.com>.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the frontmatter of the manual. In printed manuals, the document number can be found on the back cover.)

You can contact us in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library World Wide Web page:
<http://docs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system
- Send mail to the following address:
Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, California 94043-1351

SGI values your comments and will respond to them promptly.

About Networking Products

This chapter provides information about the standard hardware and software networking products provided with SGI systems. It explains the physical connection of a SGI system to an Ethernet and serial network and describes network hardware options and interface names for network devices. This chapter describes the standard networking files, directories, and daemons, and provides an overview of the network startup and shutdown processes. It also supplies a brief description of optional networking products of SGI.

Topics covered in the remaining chapters of this guide require an understanding of the fundamentals of network theory and operation. If you need information on networking fundamentals, refer to the bibliography in the introduction to this guide for additional reading. Topics in this chapter include:

- An overview of networking hardware. See “Networking Hardware” on page 1.
- An introduction to networking interface names. See “Controller Interface Names” on page 4.
- An overview of networking software. See “Networking Software” on page 4.
- A list of optional networking software products. See “Optional Networking Products” on page 5.

Networking Hardware

The networking hardware that comes standard on every SGI system is an Ethernet controller and two serial ports. (Some hardware products may have more ports than this, including an ISDN port.) The Ethernet controller may be an entire board or an integrated chip. Controllers interface between the networking software and the network medium.

If you attach the system to a network you may need additional parts. These sections address networking hardware:

- “Basic Network Attachment” on page 2
- “Networking Hardware Options” on page 3

Basic Network Attachment

To connect your Ethernet controller to a network, you must have access to an active Ethernet cable.

Figure 1-1 shows how systems (termed “stations” on the network) might be connected to an Ethernet network.




Figure 1-1 Ethernet Network Attachment

The serial ports on a SGI system allow it to connect to serial networks. Serial-line networks are systems connected by serial lines and modems. You do not need special hardware installed in your computer to connect to a serial network.

Figure 1-2 shows systems connected to a serial network using modems.




Figure 1-2 Serial Line Network

Networking Hardware Options

In addition to Ethernet and serial-line hardware, other types of controllers can be installed in SGI systems as options. Some optional hardware products are user installable, while others require installation by a System Support Engineer certified by SGI.

Optional networking products available from SGI provide support for other types of networks, including FDDI, token ring, X.25, and SNA. See your sales representative for information on the networking options available for your system.

Controller Interface Names

When you attach the system to the network, the network must know how to identify the system. The network controller, a physical board or chip, handles this function. The interface is software's interpreter and handler of the controller. The interface name is the name most evident to the user. For example, network management tools refer to the interface name when providing information about the physical controller.

To configure a controller, each network controller on a system must have a valid interface name. A single system may have multiple controllers; each controller must have a unique interface name. Several different types of controllers are available. Each type has its own special interface name. Most network software supports a maximum of four network interfaces by default.

You can get a list of the interfaces installed on a system using the *hinv* command:

```
% hinv -c network
Integral ISDN: Basic Rate Interface unit 0, revision 1.0
Integral Ethernet: ec0, version 1
```

The interface name for the Ethernet controller in this example is "ec0."

Networking Software

The standard networking software shipped with all SGI systems adheres to the Internet Model standards and protocols. It is derived from the networking software in the 4.3BSD UNIX release from the University of California at Berkeley and the RPC (remote procedure call) system from Sun Microsystems. The IRIX operating system implements the Internet Protocol suite and UNIX domain sockets using the 4.3BSD UNIX socket mechanism. The system also supports access to the underlying network media by means of raw sockets.

All standard networking software is supplied on the Execution Only Environment media (eoe and license_eoe). See Table 1-1 for a list of standard networking software for SGI systems. See Table 1-2 for a list of the optional networking products for SGI systems.

Table 1-1 Standard Networking Software

Standard Networking Software	Description
TCP/IP	Transmission Control Protocol/Internet Protocol support
UUCP	UNIX to UNIX Copy Programs
sendmail	Electronic mail support
SLIP	Serial Line Internet Protocol
PPP	Point to Point Protocol
BIND	Berkeley Internet Name Domain
FLEXlm	Flexible License Server
NCS	Network Computing System (supports NETLS only)
RPC	Remote Procedure Call support
gateway	Internet Gateway

Optional Networking Products

SGI supplies a variety of optional networking software to provide interconnectivity between various vendors and mediums. Table 1-2 lists some of these products. See your sales representative for detailed product information.

Table 1-2 Optional Networking Products

Optional Networking Software	Product Description
NFS	Includes software for Network File System (NFS); Network Information System (NIS, formerly YP); and diskless system support.
4DDN	Enables SGI systems to function as a Phase IV DECnet end node.
4DLT	Provides DECnet terminal service. (LAT)

Table 1-2 Optional Networking Products **(continued)**

Optional Networking Software	Product Description
FLEXlm License Server Developers Option	Provides FLEXlm licensing system administration tools and guidelines for integrating FLEXlm into an application.
NetVisualyzer	Offers a set of graphical traffic monitoring, diagnostic, planning, and performance analysis tools that provide network information and statistics in a visually intuitive form.
FDDI Visualyzer	Provides a graphical interface to the FDDI environment.
IRIX NetWorker	Application that automatically backs up systems over the network. Keeps online indices of all backed up files.

Planning a Network

This chapter contains common-sense approaches to planning the physical and logical aspects of your network environment. The information contained in this chapter should be read before you set up a new network or integrate into an existing network.

This chapter contains the following sections:

- “Planning the Physical Network” on page 7
- “Internet Protocol Addresses” on page 14
- “Internet Protocol Version 6” on page 18
- “Domain Names” on page 24
- “About Internet Connections” on page 25
- “About Name-to-Address Mapping” on page 32
- “Guidelines for a Subnetwork” on page 34
- “IP Address Allocation” on page 36
- “About Network Security” on page 37
- “Common Network Applications” on page 37

Planning the Physical Network

Planning the physical network requires that you first answer the question, “What network media and topology configuration would best suit the needs of my users?” A review of the MAC (Medium Access Control) level and application-level performance information about the products you are considering will help you determine the appropriate choice of media for your environment. In your review, consider the size (number of stations) of your network. Your network size will influence the media type and topology you choose for your network. If your network requires different types of

media, determine whether you have the correct equipment for integrating the various media types.

These subsections will help you answer this list of planning questions:

- What will my physical network look like? See “About Repeaters, Bridges, Routers, and Gateways” on page 8.
- Do I have a map of my network? For an example, see Figure 2-1 on page 12.
- Will I need a repeater, bridge, router, or gateway? See “About Repeaters, Bridges, Routers, and Gateways” on page 8.
- Will this network configuration meet my users’ needs? See “About Network Performance” on page 9.
- Where are my performance bottlenecks? Can I reduce or avoid them? See “About Network Performance” on page 9.

About Repeaters, Bridges, Routers, and Gateways

Your choice of media and the number of stations, networks, and protocols in your network may require the use of a repeater, bridge, router, or gateway. This section suggests the type of device required for certain network functions.

<i>repeater</i>	A device that regenerates and amplifies electrical signals. Its purpose is to extend the physical length of a network.
<i>bridge</i>	A device that decodes MAC-layer frames transmitted between different hardware and media. Its purpose is to resolve network media differences; it allows a network to be composed of various media types (Ethernet, fiber, serial, and so on). It can also be used to segment similar media types and provide segment isolation for lower network traffic.
<i>router</i>	A device that decodes and passes network-layer packets between different networks. Its purpose is to provide the physical and logical route from one network to another.
<i>gateway</i>	A device that translates protocols from one station to another. Its purpose is to allow stations with different networking protocols to communicate successfully.

Note: The terms *router* and *gateway* are sometimes used interchangeably. Be sure you know the function of the device you are considering, because the term may be technically inaccurate.

Note that each device may not be limited to a single function. For example, a gateway may also perform router functions if it is configured as a router. Table 2-1 summarizes the characteristics of each network device.

Table 2-1 Network Device Characteristics

Device Name	Media	Protocol	LAN	Purpose
repeater	same	same	same	extends physical length of the network
bridge	different/same	different/same	same/different	bridge network media differences
router	same/different	same	different	provides physical and logical route between networks
gateway	same/different	different	same/different	communication between stations with different networking protocols

About Network Performance

You can circumvent some performance bottlenecks with appropriate planning. These bottlenecks might occur as a result of your choice of media, topology, number of network devices, controller boards, or network design.

Choice of media

Be sure the capacity of the medium you have selected is adequate for the network size and data transmission type (large or small volumes of data, sporadic or steady traffic). For example, Ethernet has a range of capacities depending on the specific type of Ethernet cable used (10base5, 10base2, 10baseT or 100baseT). Media type is also a factor in data degradation. For example, 10baseT is category 3 unshielded twisted pair and is more sensitive to environmental conditions than

10base5, 100baseT must be category 5 unshielded twisted pair. These are considerations if you are planning a network for a manufacturing environment that produces a high degree of electrostatic discharge.

Number of devices

Network devices can cause degradation to the network performance. Use repeaters only when necessary to amplify the signal. Each additional device introduces additional resistance onto the network.

Choice of controller

Choose the most efficient controller for your media. For example, Silicon Graphics supplies a standard Ethernet controller. An optional Efast™ card handles more of the protocol processing in hardware and frees the station's CPU for other processing.

Design of network

Think about the design of your network before you begin setting it up. If possible, put departments that interact heavily on the same network to decrease router traffic. Use dedicated routers to handle heavy traffic between networks.

About Wide Area Networks

In addition to the many options available for constructing local area networks, there are several different ways of connecting local area networks into wide area networks. These systems can be used to tie together local area networks at different locations, to allow users working at scattered locations to access a network, and to connect your network to the outside world. These subsections introduce the different systems:

- “Serial Line Internet Protocol (SLIP)” on page 12
- “Point to Point Protocol (PPP)” on page 13
- “UNIX to UNIX Copy Program (UUCP)” on page 13
- “Integrated Services Digital Network (ISDN)” on page 13
- “Internet Gateway” on page 13
- “High Performance Wide Area Networks” on page 14

Figure 2-1 shows how different kinds of wide-area connections might fit into a large heterogeneous network.

Two of the available systems, Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP), provide a way of transferring Internet Protocol (IP) packets over a serial telephone line; this means that SLIP and PPP users can access network resources much as if they were on the local area network. PPP can also be used with Integrated Services Digital Network (ISDN), available on certain platforms. ISDN uses a high-speed digital telephone line to achieve higher throughput than is possible with a modem connection.

Another system, UNIX to UNIX Copy Program (UUCP), is an older system, primarily designed for transferring information (such as network news and electronic mail) in batch mode over serial lines.

Higher-performance network connections can be made using specialized hardware. These connections are usually over dedicated lines, leased from a telephone company, or over the telephone company's packet-switched network.




Figure 2-1 Heterogeneous Network With Wide-Area Connections

Serial Line Internet Protocol (SLIP)

SLIP provides simultaneous operation of multiple processes on a serial cable or telephone line. It allows network users the freedom to use TCP/IP based applications over a serial cable or modem connection.

You might consider setting up a SLIP network when cost and distance are large factors in your network planning.

Point to Point Protocol (PPP)

The Point to Point Protocol is similar in nature to SLIP. PPP provides a network connection as if your system were connected to the remote host by a LAN connection. Multiple processes and TCP/IP based applications are supported.

UNIX to UNIX Copy Program (UUCP)

UUCP, also called the Basic Networking Utilities (BNU), is a set of utilities that lets stations using a version of the UNIX operating system (such as IRIX) communicate with each other over serial lines. The utilities provided range from those used to copy files between computers to those used for remote login and command execution.

You may consider setting up UUCP for long-haul communications using modems and telephone lines. It is usually used to distribute electronic mail and network news.

Integrated Services Digital Network (ISDN)

ISDN is a system that connects systems using high-speed digital telephone lines. ISDN can achieve throughput up to 128 kilobits (Kb) per second, several times faster than normal modem connections. However, ISDN service can be expensive, and is not available in all areas, nor all platforms. See the *ISDN User's Guide* for more information on ISDN.

Internet Gateway

The Internet Gateway provides a server process to connect with the Internet and a means of configuring various name services. It acts as a router, and has its own system of help screens to help the user through the configuration process.

High Performance Wide Area Networks

If you need higher performance than you can get using SLIP or PPP over a modem link or ISDN, there are several choices available to you. The choices include Frame Relay networking and leased line service, from 56 Kb (56 Kb per second) to T1 (1.5 Mb per second) and T3 (up to 45 Mb per second). If you require this kind of service, you will have to shop around, comparing the prices and services offered by local Internet service providers.

Internet Protocol Addresses

Each system on your network needs a unique Internet Protocol (IP) address for each of its network interfaces. The Internet Network Information Center (InterNIC) is responsible for assigning the network portion of an Internet address for each site. For example, if Company A applies for an Internet address, the InterNIC provides the network portion of the Internet address for the entire Company A site. A centralized organization within Company A is responsible for assigning and managing the station ID portion of the Internet address.

Internet addresses are maintained on each station or in a centralized network database such as NIS or BIND. See “About Name-to-Address Mapping” on page 32 for a comparison of the different database types. Each station that wishes to communicate must have a valid Internet address registered in the appropriate database. The standard hosts name-address database on IRIX stations is the */etc/hosts* file.

These subsections will help you answer this list of planning questions:

- What is an IP address? See “Format of Internet Protocol (IP) Version 4 Addresses” on page 15.
- How do I obtain a valid Internet address for my site? See “Obtaining a Network Number” on page 17.
- What information do I need to gather before I can obtain an Internet address? See “Required Information for Obtaining an Internet Address” on page 18.

Format of Internet Protocol (IP) Version 4 Addresses

An IP address is a 32-bit number that network software uses to identify a system on a network. For the sake of human readability, these addresses are usually represented as four one-byte integers, separated by dots (for example, 150.166.248.17). Every system on an IP network must have its own unique IP address for the network to function properly. Systems with more than one network interface must have a unique IP address for each interface.

Note: Unlike a system's Ethernet address, a system's IP address is determined by the network and network system administrators.

Conceptually, each 32-bit IP address is a pair of numbers where one number represents the network and the other the system itself. There are four classes of addresses in use (A through D). The class of address is determined by the first bits of the address:

- Class A addresses begin with 0 and have 7 bits for the network number and 24 bits for the host number.
- Class B addresses begin with 10 and have 14 bits for the network number and 16 bits for the host number.
- Class C addresses begin with 110 and have 21 bits for the network number and 8 bits for the host number.
- Class D addresses begin with 1110 and are special "multicast" addresses for use within a network site.

In all cases, host numbers 0 and 255 are reserved, and may not be used for actual systems.

Figure 2-2 shows the format of the different classes of Internet addresses.

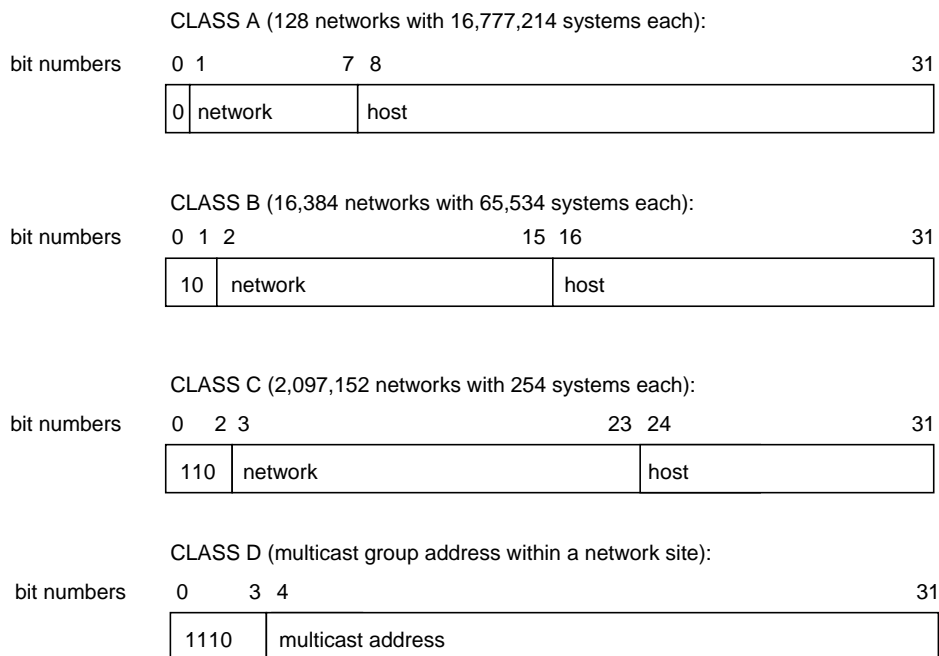


Figure 2-2 Format of Internet Protocol (IP) Version 4 Addresses

To simplify Internet addressing, dotted decimal notation is used to break the 32-bit number into four decimal numbers separated by dots.

For example, the IP address 128.74.41.123 in binary is

```
10000000 | 01001010 | 00101001 | 01111011
```

or

```
128 | 74 | 41 | 123
```

Class A, B, and C IP addresses in dot notation conform to the following specifications:

Class A -- 001.*hhh.hhh.hhh* through 126.*hhh.hhh.hhh*

Class B -- 128.001.*hhh.hhh* through 191.254.*hhh.hhh*

Class C -- 192.000.001.*hhh* through 223.255.254.*hhh*

Note: *hhh* is the local system and the leading numbers are the network.

Networks are usually identified by *network numbers*—IP addresses in which the host portion is not specified. For example, 150.166 represents a Class B network, and 192.26.80 represents a Class C network.

If your network will be connected to the Internet, then you must obtain a unique network number, as described in “Obtaining a Network Number” on page 17. All the systems on your network must have IP addresses allocated from your network.

If you are adding a machine to an existing network, its IP address must be allocated from that network.

Obtaining a Network Number

You should obtain an Internet network number before you begin setting up your network. The allocation of network numbers is managed by a set of organizations called Network Information Centers (NICs). (See “About Local Network Information Centers” on page 27.) If your network is going to be isolated, and will *never* be attached to the Internet, you can theoretically use any addresses you like. However, if your network is ever going to be attached to the Internet, you should obtain a valid network number. Before you request the network number, you should determine the current needs of your organization (how many systems do you currently have that should be on the network?) and expected growth over the next five years.

There are several ways to obtain a network number. In many cases the best option, if you are connecting to the Internet through an Internet service provider, is to have the service provider assign you a portion of the address space they have been allocated by the local NIC.

The InterNIC recommends that you request a network number from your network service provider. If they cannot supply one, contact your provider’s provider. As a last resort, contact your Network Information Center. See “About Local Network Information Centers” on page 27.

Required Information for Obtaining an Internet Address

To request an Internet network address, you typically need to supply the following information to the local NIC:

- Your administrative point of contact (POC). The administrative POC is the person responsible for answering administrative and policy questions about the network. You need to know his/her name, title, mailing address, and phone number.
- Your technical point of contact (POC). The technical POC is responsible for the technical support of the network. You need to know his/her name, title, mailing address, and phone number.
- Organization name and postal address.
- Your network name (up to 12 characters).
- Your network's geographic location and organization name.
- The name and location of the network document plan.
- Gateway information (connectivity, hardware, software, address).
- The approximate size of your network (number of hosts and subnets), initially and within one year.
- Type of network (research, educational, government non-defense, commercial).

If you already have one or more network numbers assigned to your organization, the NIC may require you to provide information on how these are being used, as evidence that you really need a new network number.

If you request 16 or more Class C network numbers, the InterNIC requires you to provide information on network topology, and if you request 256 or more Class C network numbers or a Class B network number, the InterNIC requires you to provide a diagram of the proposed network.

Internet Protocol Version 6

This section describes Internet Protocol Version 6 (IPv6) and covers the following topics:

- "IPv6 Overview" on page 19
- "IPv6 Addressing" on page 20
- "Format of Internet Protocol IP Version 6 Addresses" on page 21

- “IPv6 Header Format” on page 22
- “Installing and Configuring IPv6” on page 24

IPv6 Overview

Internet Protocol Version 6 (IPv6) is the next generation Internet protocol designed to eventually replace the current IPv4 (see “Internet Protocol Addresses” on page 14 and “Format of Internet Protocol (IP) Version 4 Addresses” on page 15).

IPv4 uses a 32-bit address scheme to represent a unique Internet address while IPv6 uses a 128-bit addressing scheme. Because IPv6 has four times the number of bits to represent addresses, it provides many billion times the address space of IPv4 and adds an almost unlimited number of available IP addresses. This resolves a projected future shortage of IPv4 addresses as more and more devices are added to the Internet.

The IPv6 address scheme provides stateless, automatic configuration; improved service capabilities for services such as real-time; and enhanced support for authentication and privacy capabilities. It can help to restore end-to-end architecture in networks. To ease the transition from IPv4-based networks to IPv6-based networks, IRIX supports a dual-stack interface that supports both IPv4 and IPv6 protocols on the same system. The IPv6 basic socket API supports applications that transparently use the correct Internet Protocol. Applications can choose the right protocol to communicate with any particular node.

IPv6 has an encapsulation mechanism called tunneling that allows IPv6 packets to operate over IPv4 networks. The IPv6 protocol allows networking applications to operate transparently in both environments.

IRIX supports IPv6 using Berkeley sockets system calls. The differences between using sockets for IPv4 networking and IPv6 networking are minimal. The UDP and TCP transport layer protocols work in exactly the same way over IPv6 as they do over IPv4. For more information on IPv6 socket creation and socket options, see the `ipv6(7P)` man page.

The *IRIX Network Programming Guide* also includes information about IPv6 socket creation and socket options.

For a list of Internet Engineering Task Force (IETF) request for comments (RFCs) documents that describe the IPv6 functionality supported by IRIX, see the `ipv6(7P)` man page.

IPv6 Addressing

IPv6 supports unicast and multicast addressing. Descriptions are as follows:

- unicast
An IPv6 address for a single interface.
- multicast
An identifier for a set of interfaces that typically belong to different nodes.

IPv6 has different types of unicast addresses. A *link-local* unicast address is used over a single link. A *global* unicast address is used across a global Internet.

The `ndpd(1M)` daemon is invoked at boot time to perform IPv6 neighbor discovery for the automatic configuration of link-local addresses and to automatically configure global addresses when router advertisements are received from an IPv6 aware router.

The neighbor discovery protocol (NDP) is used to dynamically map between IPv6 Internet and device addresses such as Ethernet addresses. It is used by all the Ethernet interface drivers, although it is not specific to Ethernet.

The `ndpd` daemon is normally started by the network start-up script during boot and is run with the options contained in `/etc/config/ndpd.options`. When it starts, it first starts the NDP on each interface by choosing a link-local address and sending a neighbor solicitation message to check that the address is unique for that interface. It also sends out router solicitation messages on each interface. When a router advertisement message is received, it will try and configure a global address on that interface using the prefix contained in the message.

The `route6d` daemon is a routing daemon that supports routing information protocol (RIP) requests over IPv6.

Additional information on IPv6 topics described in this overview are available as follows:

- For more information on the `ndpd` daemon, see the `ndpd(1M)` man page.

- For more information on neighbor discovery protocol (NDP), see the `ndp(7P)` man page.
- For more information on the `route6d` daemon, see the `route6d(1M)` man page.
- For detailed information on NDP, see RFC 2461, *Neighbor Discovery for IP Version 6 (IPv6)*.
- For detailed information on IPv6 addressing, see RFC 2373, *IPv6 Addressing Architecture*.

Format of Internet Protocol IP Version 6 Addresses

IPv6 uses a 128-bit number to represent an IP address as shown in Figure 2-3.




Figure 2-3 Format of Internet Protocol (IP) Version 6 Addresses

IPv6 addresses are classless.

IPv6 addresses can be represented as text strings as follows:

- `x:x:x:x:x:x:x:x`

Each `x` represents a hexadecimal value of a 16-bit piece of the address (see Figure 2-3).

An example address follows:

```
1080:0:0:8:800:200C:417A
```

It is not necessary to write the leading zero in an individual field, but there must be at least one numeral in every field except when zeros are compressed, as in the next example.

- `x::x:x:x:x`

You can use the single double colon `::` syntax to indicate multiple groups of 16-bits of zeros. The `::` syntax can appear only once in an address. You can also use this double colon syntax to compress the leading or trailing zeros in an address.

An example follows:

`1080::8:800:200C:417A` represents the same address as
`1080:0:0:8:800:200C:417A`

- `x:x:x:x:x:x:d.d.d.d`

Each `x` represents the hexadecimal values of the six, high-order, 16-bit pieces of the address. Each `d` represents the decimal values of the four, low-order, 8-bit pieces of the address (standard IPv4 representation).

An example follows:

`0:0:0:FFFF:129.144.52.38`

In IPv6, the left, or most significant bits, of the IP number identify the network on which a machine resides and are the *prefix* bits. The right, or least significant bits, identify a machine residing on a network or subnetwork and are the *interface identifier* bits. The prefix and interface identifier bits are shown in Figure 2-4.




Figure 2-4 IPv6 Prefix Bits and Interface Identifier Bits

IPv6 Header Format

Figure 2-5 shows the IPv6 header format specified by IPv6 protocol. The IPv6 headers are all the same length and contain fewer fields than IPv4. IPv6 headers are terminated on 64-bit boundaries rather than 32-bit boundaries as in IPv4.

Version	Traffic class	Flow label	
Payload length		Next header	Hop limit
Source IP address			
Destination IP address			
Data portion of datagram			

Figure 2-5 IPv6 Header Format and Data Portion of a Datagram

The IPv6 header fields are briefly described as follows:

- In IPv6, the *version* field must equal 6.
- The semantics of *traffic class* and *flow label* are not yet agreed upon and these fields have default values of zero.
- *Payload length* defines the length of the packet payload.
- *Next header* refers to a higher-layer protocol, such as TCP, or may indicate an IPv6 extension header.
- *Hop limit* contains a value that is decremented as nodes forward the packet.
- *Source IP* address is the 128-bit address of the originating node.
- *Destination IP* address is the 128-bit address of the target recipient for the packet.

IPv6 supports optional, Internet-layer information that is encoded in separate headers called *extension headers*. Extension headers may be placed between the IPv6 header and the upper-layer header in a packet. There are a small number of such extension headers, each identified by a distinct next header value. For more information on extension headers, see *RFC 1883* and *RFC 2460*.

Installing and Configuring IPv6

To enable IPv6 protocol on a system, perform the following:

1. Use the `inst(1M)` utility to install the `oe.sw.ipv6` from your IRIX distribution media. The `oe.man.ipv6` and `oe.man.ipv6_rn` subsystems are optional documentation subsystems.
2. Enable IPv6 protocol within the kernel by using the `sysune(1M)` command to set `ip6_enable` as follows:

```
sysune ip6_enable 1
```

If your system is a router, you need to enter the following command:

```
sysune ip6forwarding 1
```

If your system is an IPv6 router, you also need to modify the arguments to `ndpd` and `route6d`. For information on how to do this, see the `ndpd(1M)` and `route6d(1M)` man pages.

If you want to process routing headers, enter the following command:

```
sysune ip6forwsrcrt 1
```

Domain Names

If you are planning on putting your site on the Internet or exchanging e-mail with sites on the Internet, you should register a domain name with your local NIC. A domain name uniquely identifies your organization. For example, Silicon Graphics has the domain name `sgi.com`.

These subsections explain domain names and subdomains:

- How do I register a domain? See “Obtaining a Domain Name” on page 25.
- If you need to subdivide your domain, see “About Subdomains” on page 25.

The Internet uses Domain Name Service (DNS) to map domain names to IP addresses. Therefore, even if you do not use DNS internally, you must provide DNS name servers on the Internet in order to connect your network to the Internet. You should have at least two name servers, a primary and a secondary server. For robustness, the secondary server should not be connected to the Internet through the same gateway as the primary

server. Since many organizations are not big enough to have multiple gateways to the Internet, a common solution is to make a reciprocal arrangement with another organization to provide secondary name service for each other.

If you are connecting to the Internet through an Internet service provider, they may be able to provide name service for your organization, or help you locate someone to provide secondary name service if you are able to provide a primary name server.

Obtaining a Domain Name

As with network numbers, the registration of domain names is administered by the Network Information Centers. In some cases, there is a fee associated with holding a domain name. For example, the InterNIC currently charges a fee of \$100 for the first two years, and \$50 a year thereafter for domains under its jurisdiction.

You can register your domain name through your local NIC. See “About Local Network Information Centers” on page 27 for contact information. When you register a domain, you should also register a *reverse domain*, also known as an *IN-ADDR* domain. The reverse domain provides a mapping from IP addresses to domain names.

In many cases, Internet service providers will register your domain for you, for a fee.

About Subdomains

Once you have a domain name registered, you are free to establish subdomains of your own. This is particularly useful for large organizations that use the Domain Name Service (DNS). The use of subdomains with DNS allows some administrative chores to be decentralized.

For example, suppose salad.com has branch offices in Gilroy and Paris. These could be established as subdomains, gilroy.salad.com, and paris.salad.com.

About Internet Connections

Chances are, you will want to connect your system or network to the Internet. Wherever you may be, there is likely an Internet gateway available within your local calling range. The following sections offer some information that should help you get set up and

running. Obviously, each situation is somewhat different, and your local service provider will have variations in service and equipment. Some research and experimentation is usually required before everything works smoothly.

- How much of the Internet will I need? See “Before Connecting to the Internet” on page 26.
- How do I contact the Internet? See “About Local Network Information Centers” on page 27.
- What can I get from the Internet? See “Online Information Sources” on page 30.

Before Connecting to the Internet

Before you sign up for an internet connection, consider what level of service you need. For example, if you are an individual looking for basic e-mail, news, and file transfer capabilities, it probably would not make sense to install a dedicated network cable in your home for economic reasons. A better choice for single-user access might be to subscribe to a network provider who establishes an account for you on their system (one that is currently connected to the Internet). Typically, access to their system is through a modem connection.

If you are trying to establish a connection to the Internet for a corporation, you will likely need the bandwidth of a leased line, and all the required hardware that goes with it. You will have to take into consideration the many administrative issues of running a site. These issues include, but are not limited to

- Cost analyses/budget
- Establishing a domain
- Applying for IP addresses
- Establishing site policy
- Establishing site security
- Administration of network services (such as Domain Name Services, NIS, e-mail, and so on)

There are providers of network connectivity that can provide varying levels of service. You must investigate the providers, and decide who provides the level of service you need, at the appropriate cost.

If you choose an individual account on a provider's machine, the service provider deals with most, if not all, of the administrative tasks, and you simply enjoy access to the Internet.

If you would like a broader range of services, most providers will set you up with a dedicated modem and phone line for your exclusive use, or they can provide a network-only service (using SLIP, PPP, or UUCP), either through modems or other network connections.

Connecting your network to the Internet requires a number of steps, including arranging name servers, obtaining a network number, and registering a domain name for your organization. Many Internet service providers are willing to provide these services for a fee.

If you are trying to set up internet access for a company, or corporation, you should research the issues listed above. Based on the information you obtain, formulate a plan for your site based on the needs and expectations of your organization. One of the best sources of information is the Internet itself. You should first obtain an individual account from a local provider. With the individual account, you can gain access to a large amount of information pertaining to establishing a site on the Internet.

About Local Network Information Centers

Before you connect your site to the Internet, you will need to contact your local Network Information Center. The assignment of network numbers and domain names is coordinated by the Network Information Centers. There are three main regional Network Information Centers, as shown in Table 2-2.

Table 2-2 Network Information Centers

Region	Organization
Asia/Pacific	Asia Pacific Network Information Center (APNIC)
Europe	Réseaux IP Européens Network Coordination Centre (RIPE NCC)
Americas	Internet Network Information Center (InterNIC)
other areas	InterNIC

Procedures for obtaining IP addresses and registering domain names vary, so contact your local NIC for information.

Internet Network Information Center

The Internet Network Information Center (InterNIC) was formerly the sole Network Information Center. It serves as the primary NIC for most of North and South America, as well as for other regions that do not yet have NICs of their own. InterNIC maintains a large archive of informational documents, which can be accessed using WWW, FTP, or by e-mail to an automated-response mail server. Registration authority for some countries (including Canada and Brazil) is delegated to national NICs. Contact information for the national NICs may be obtained through InterNIC.

Network Solutions
Attn: InterNIC Registration Services
505 Huntmar Park Drive
Herndon, VA 22070
Phone: 1-800-444-4345 or 1-703-742-4777
E-mail: question@internic.net (general inquiries)
E-mail: hostmaster@internic.net (registration services)
WWW: <http://www.internic.net/>
FTP: <ftp.ds.internic.net> (complete RFCs, and so on)
FTP: [rs.internic.net](ftp.rs.internic.net) (registration information)
E-mail server: mailserv@rs.internic.net (send message with subject line "HELP")

Réseaux IP Européens

Réseaux IP Européens (RIPE) maintains an NIC that provides registration services for European sites. It also maintains a store of informational documents, including the InterNIC's FYI documents, and instructions on how to register a host or network in one of the European domains.

RIPE Network Coordination Centre
Kruislaan 409
NL-1098 SJ Amsterdam
The Netherlands
Phone: +31 20 592 5065
Fax: +31 20 592 5090
E-mail: ncc@ripe.net
WWW: <http://www.ripe.net/>
FTP: <ftp.ripe.net>

Asia Pacific Network Information Center

The Asia Pacific Network Information Center (APNIC) coordinates network information for the Asia and Pacific region. Registration authority for some countries is delegated to national NICs. Contact information for the national NICs may be obtained through APNIC.

Asia Pacific Network Information Center
c/o United Nations University
53-70 Jingumae 5-chome
Shibuya-ku, Tokyo 150
Japan
Phone: +81-3-5467-7014
Fax: +81-3-5276-6239
E-mail: info@apnic.net
WWW: <http://www.apnic.net/>
FTP: [archive.apnic.net](ftp://archive.apnic.net)

Online Information Sources

With an individual account or other access to the Internet, you can get the information you need to provide access to your own site.

Usually, the provider of an individual account will also provide new-user documentation that describes the basics of using the Internet. You can use the World Wide Web (WWW) and the File Transfer Protocol (FTP) to access a wealth of information on many subjects, including Internet connectivity. If you do not know how to use FTP, see “Retrieving Files With Anonymous FTP” on page 31 for a short tutorial. How you access the Web depends on what Web browser you are using. Most Web browsers have online help available.

The following subsections will help you use the Internet:

- Can the Network Information Center help me connect to the Internet? See “Network Information Centers” on page 30.
- How can I find an Internet provider? See “Internet Society” on page 30.
- When I get there, how do I get the files onto my system? See “Retrieving Files With Anonymous FTP” on page 31.

Network Information Centers

Your local NIC maintains archives of useful information on connecting to the Internet. In addition to information about requesting network numbers and registering domain names, they may have lists of local service providers. Most NICs make this information available by WWW and FTP. See “About Local Network Information Centers” on page 27 for WWW and FTP addresses for the major NICs.

The InterNIC has produced a series of information bulletins called FYIs. Especially notable is FYI 16, entitled *Connecting to the Internet—What Connecting Institutions Should Anticipate*. While this is aimed primarily at U.S. educational institutions, it remains one of the better pieces of documentation on establishing a site on the Internet. The FYI documents are available by WWW and FTP from the InterNIC and from RIPE.

Internet Society

The Internet Society is a non-governmental international organization for global cooperation and coordination of the Internet. They also provide useful online

information—in particular, information on finding an Internet service provider, and a list of network service providers around the world. This information is available by WWW. A subset is available by anonymous FTP.

WWW: <http://www.isoc.org/>

FTP: <ftp.isoc.org>

Retrieving Files With Anonymous FTP

Anonymous FTP is a conventional way of allowing you to sign onto a computer on the Internet in order to obtain copies of files that are made available to the public. Some sites offer anonymous FTP accounts to distribute software and various kinds of information. If you have never used *ftp*, here is a brief summary on how to use the *ftp* command. To connect to a remote host, specify the hostname on the command line:

```
ftp ftp.ds.internic.net
```

When *ftp* connects with the remote system, it prompts you for a login name. Use the login name “anonymous”:

```
Connected to ftp.ds.internic.net.  
Name (ftp.ds.internic.net:guest): anonymous  
331 Guest login ok, send ident as password.  
Password:
```

Many systems allow any password and request that the password you choose is your user ID. If this fails, the generic password is usually “guest.”

```
230 Guest login ok, access restrictions apply.  
Remote system type is UNIX.  
Using binary mode to transfer files.  
ftp>
```

Once connected and logged in, you can use *ftp*'s *cd* and *ls* commands to look at the files available on the remote system. To obtain a file from the remote system, use the *get* command. The *get* command copies one file from the remote system to your local system. To obtain multiple files from the remote system, use the *mget* command.

```
ftp> cd fyi  
250 CWD command successful.  
ftp> get fyi6.txt  
local: fyi6.txt remote: fyi6.txt  
200 PORT command successful.  
150 Opening BINARY mode data connection for fyi6.txt (3459 bytes).
```

```
226 Transfer complete.  
3459 bytes received in 0.46 seconds (7.34 Kbytes/s)  
ftp>
```

About Name-to-Address Mapping

Because IP addresses are difficult to remember, they are usually associated with names. In the case of a machine with a single IP address, this name *usually* consists of the machine's hostname and domain name. For example, a machine called fruit in the domain salad.com would usually be referred to as fruit.salad.com. For clarity, this type of name will be referred to in this section as a *network connection name*.

Because network connection names usually correspond to the machine's hostname, these network connection names are commonly referred to as "hostnames," but this can be misleading. The actual hostname is defined in the `/etc/sys_id` file. By default, this hostname is used as the network connection name for the machine's primary network interface, but this behavior is configurable. A machine with multiple network interfaces has multiple network connection names associated with it. By convention, each of these connection names contains the hostname—for example, if the host fruit acts as a gateway between two networks in the salad.com domain, it might use these names:

```
fruit.salad.com  
gate-fruit.salad.com
```

The process of mapping network connection names to IP addresses is commonly called *hostname resolution*. There are several different systems for hostname resolution. Machines can use a local database (the `/etc/hosts` database), or they can obtain information from servers on the network, using either the Network Information Service (NIS) or the Domain Name System (DNS). The following sections describe the advantages and drawbacks of the different systems:

- “`/etc/hosts` Database” on page 33
- “Domain Name System” on page 33
- “Network Information Service (NIS)” on page 33

/etc/hosts Database

The */etc/hosts* database is an ASCII file that you can modify with any text editor. The file contains lines of text that specify IP addresses and network connection names.

For a small network of stations under the same administrative control, maintaining a consistent */etc/hosts* database is straightforward. Establish a master copy on one station and make additions or deletions from its file. Then use *rcp* or *rdist* to copy the file to the other stations in the network.

Maintaining consistent versions of */etc/hosts* on every station in a large network is troublesome. NIS and the BIND name server both make maintenance easier by providing a centralized version of the host database.

Domain Name System

The Internet uses the Domain Name System (DNS) to map names to IP addresses. The most common implementation of a DNS name server is called Berkeley Internet Name Domain (BIND). If your network interfaces with the Internet, you must have at least two DNS name servers, a primary and a secondary server. Your Internet service provider may be able to take care of this requirement for you.

BIND is best suited for large networks, or networks connected directly or indirectly to the Internet. BIND provides access to a much larger set of stations than is provided in the */etc/hosts* database. A drawback of BIND is its complicated setup. BIND is described in more detail in Chapter 6, "BIND Name Server."

Network Information Service (NIS)

NIS is a network-based information service and an administrative tool. It allows centralized database administration and a distributed lookup service. NIS supports multiple databases based on regular text files. For example, NIS databases can be generated from the *hosts*, *passwd*, *group*, and *aliases* files on the NIS master.

NIS is best suited for a moderate-sized network (one containing approximately 1000 stations, or a small collection of interconnected networks). NIS is part of the NFS optional software and is detailed in the *NIS Administration Guide*.

Guidelines for a Subnetwork

Subnetting allows you to divide a single network into a set of subnetworks. Subnetworks are useful for many reasons. For example, if you have a satellite office that connects to your main network, it should have its own network number or subnet. If you have a large number of systems to be connected by Ethernet, you may have to use subnets to overcome physical limitations on the number of hosts and length of network cable that can be supported on a single Ethernet network.

Subnetting should be considered when the class limits are unrealistic for your network. For example, a Class B network gives you approximately 64,000 stations per network. This far exceeds the maximum number of stations allowed on most networks. Subnetting allows the local organization to designate some of the host ID bits to form a subnet. Subnetting generates a realistic number of stations per network. All changes are made at the local site by the site administration group and are transparent to off-site stations.

Planning is required for subnetting a network (see “Subnetting a Network” on page 55 for subnetting procedure). Primarily, you must determine how to partition the host part of the 32-bit Internet address. To define local subnetworks, use bits from the host number sequence to extend the network portion of the Internet address. This reinterpretation of IP addresses is done only for local networks. It is not visible to off-site stations. You should have at least a rough idea of the physical layout of the network before you plan your subnets. For example, you might want to have a subnet for each floor of your building. If you have a branch office that’s connected to your main network, you might want to set aside one or more subnets for it. In some cases, you may want to set aside a subnet for SLIP and PPP clients (see “SLIP and PPP Routing and Address Allocation” on page 146).

Sites with a Class A network number have 24 bits of host part with which to work; sites with a Class B network number, 16 bits; and sites with a Class C network number, 8 bits. For example, if your site has a Class B network number, each station on the network has an Internet address that contains 16 bits for the network number and 16 bits for the host number. To define 254 local subnetworks, each possessing at most 254 stations, you can use 8 bits from the host portion of the address. Construct new network numbers by concatenating the original 16-bit network number with the extra 8 bits containing the local subnetwork number.

Note: It is highly recommended that the size of the portion of the host number sequence used for a subnet id be the same for all subnets. Variable subnets are supported but difficult to get correct.

Figure 2-6 shows what happens to the bit assignments in a Class B Internet address that is subnetted.

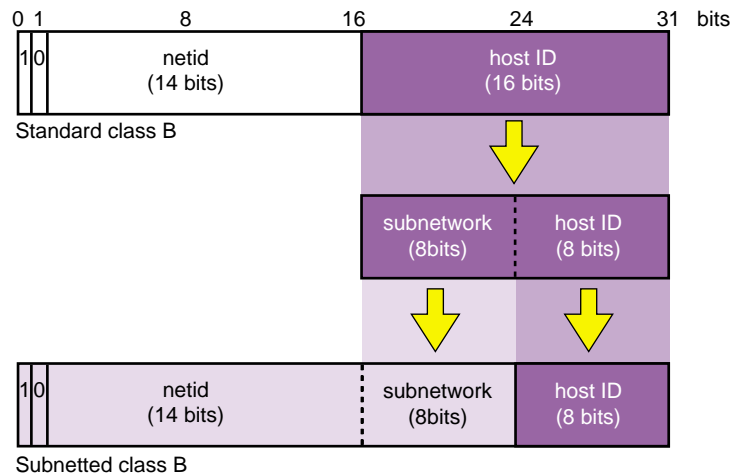


Figure 2-6 Subnetted Class B Address

For example, the Class B Internet address for an entire site as seen from other sites might be 128.50. If subnetting is enabled within the site, the site might be composed of several subnets with network IDs like 128.50.20, 128.50.21, 128.50.22, and so on. A station that resides on the subnet 128.50.21 might have the Internet address 128.50.21.5.

Note: Numbers consisting of all 0s, and all 1s, are reserved for broadcast addresses. Do not use subnetwork numbers with all 0s or all 1s.

IP Address Allocation

Once you have obtained a network number, decided which system to use for hostname resolution, and decided whether you are going to subnet the network, you are ready to allocate IP addresses for individual systems. For most systems, this is as simple as assigning an unused IP address from the correct net or subnet. If you use the syntax of the */etc/hosts* file, the result will look something like this:

```
150.26.80.1    green.salad.com green
150.26.80.2    tossed.salad.com tossed
150.26.80.3    jello.salad.com jello
<IP address>  <host>.<domain> <host> <alias>
```

Note: Host numbers 0 and 255 are reserved, and shouldn't be used.

Systems with more than one network interface may be connected to more than one subnet, and require one address for each connected interface. Each interface should be assigned an address from the subnet that the interface is connected to. For example, if *fruit.salad.com* acts as a gateway between the 150.26.80 net and the 150.26.42 net, it might have the following entries:

```
150.26.80.19   fruit.salad.com fruit
150.26.42.1    gate-fruit.salad.com gate-fruit
```

Even if you are planning on using NIS or BIND for hostname resolution, you will probably want to put together an */etc/hosts* file. If you install this on your systems as you attach them to the network, you will be able to communicate while you get NIS or BIND up and running.

You should also establish some policy for allocating IP addresses for new systems once the network is in place. If your organization is large, you might want to delegate this authority to separate organizational units. For example, the branch office with its own subnet should allocate IP addresses as needed from its subnet. If your organization is divided up into subdomains, you might want to assign authority over certain subnets to subdomain administrators.

About Network Security

Securing a network is difficult. If you can discourage potential intruders and quickly isolate or pinpoint successful intruders, you can consider your network secure. You should establish a plan for keeping your network secure before you connect your network to the Internet. For information on network security, see Chapter 5, “Network Security,” in *IRIX Admin: Backup, Security, and Accounting*.

Common Network Applications

This guide is written specifically to support the standard network hardware and software—Internet protocols over Ethernet. However, when discussing networking in general, it is difficult to ignore network applications that are not standard, but are common to most network environments.

These subsections present a brief overview of some of the common network applications that you should consider when planning your network:

- For quick word about mail, see “About Electronic Mail” on page 37.
- For a brief look at remote file sharing, see “About Network File System (NFS)” on page 37.

About Electronic Mail

Electronic mail is a group of programs (*sendmail*) used to send and receive messages to and from users on the same local station or between remote stations. Mail can be sent using UUCP or TCP/IP protocols. IRIX supports both System V (*/bin/mail*) and 4.3BSD (*/usr/sbin/Mail*) mail programs, as well as most other mailers including *Netscape Mail*, which provides a graphical interface for electronic mail.

About Network File System (NFS)

NFS is a network program that can access a remote station’s filesystem and attach it and its data to the local station’s filesystem. On the local station, the remote filesystem is accessed as if it were local.

NFS should be considered in a network when you want to share files between stations. With NFS, software or data used by a group is put on an NFS server. Authorized NFS clients access the data over the network when needed. This approach ensures consistent information, frees up disk space on client stations, and simplifies the backup procedure. NFS is an optional software product and is described in the *ONC3/NFS Administrator's Guide*.

Note: NFS is not included with the IRIX operating system and must be purchased separately.

Setting Up a Network

This chapter contains the following sections:

- “Configuring a System for a Network” on page 39
- “Setting Up a Router” on page 47
- “Subnetting a Network” on page 55
- “Modifying the Network Interface Configuration With the `/etc/config/netif.options` File” on page 56
- “Changing Network Parameters in the `ifconfig-#.options` File” on page 62
- “Configuring `/etc/gateways` Files for Networks That Do Not Support Broadcast or Multicast” on page 64
- “Configuring Multiple Network Interfaces” on page 67
- “Dynamic Host Configuration With Proclaim” on page 68
- “Creating a Local Network Script” on page 91
- “Turning On Remote Access Logging” on page 91
- “Setting Up Network-Wide Services” on page 92
- “Setting up Time Synchronization Service” on page 103
- “Reserving Resources With RSVP” on page 105
- “Troubleshooting Your Ethernet Connection” on page 108
- “Configuring a Dual-port Gigabit Ethernet Board” on page 112

Configuring a System for a Network

The procedure in this section explains how to configure a workstation, with one interface, for an Ethernet network using its local `/etc/hosts` file (without BIND or NIS). Configuring the station takes six steps:

1. Bring the station down.
2. Attach the station to the network.
3. Check the station's network configuration.
4. Modify the */etc/hosts* database.
5. Name the station.
6. Test the connection.

Each of these steps is explained in the sections that follow:

- “Attaching Your Station to an Ethernet Network” on page 40
- “Checking Your Ethernet Connection” on page 40
- “Checking the Network Software Configuration” on page 43
- “About the Hosts Database” on page 43
- “Modifying the Hosts Database” on page 45
- “Naming Your Station” on page 46
- “Testing Your Network Connectivity” on page 46

Attaching Your Station to an Ethernet Network

Attach your station to the network by connecting one end of the Ethernet cable to the I/O port on the back of your station, and the other end into the network.

Checking Your Ethernet Connection

You can use the *ping* command to check your Ethernet connection. This command tests whether you can connect with another system on the Ethernet network. Perform the following steps:

1. Obtain the hostname of at least one reliable station on the local area network to which your system is connected. If possible, get the fully qualified hostname and the IP address. (For example, a hostname might be hancock, and the fully qualified hostname might be hancock.corp.gen.com, while the IP address might be 128.70.3.56.) It is important that the station you select has a reliable Ethernet connection and that it is up and running.

2. Once you have obtained a hostname and IP address, enter the command

```
ping -r hostname
```

You should see a series of records indicating the returned packets from the remote host. For example (using our example system):

```
PING hancock (128.70.3.56): 56 data bytes
64 bytes from 128.70.3.56: icmp_seq=0 ttl=255 time=2 ms
64 bytes from 128.70.3.56: icmp_seq=1 ttl=255 time=2 ms
64 bytes from 128.70.3.56: icmp_seq=2 ttl=255 time=2 ms
64 bytes from 128.70.3.56: icmp_seq=3 ttl=255 time=2 ms
64 bytes from 128.70.3.56: icmp_seq=4 ttl=255 time=2 ms
```

3. Press Ctrl+C or your Delete key to stop the *ping* command. You see the tallied results of the *ping* command. For example:

```
---- Hancock PING Statistics ----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 2/2/2 ms
```

4. If your network connection is working, you should see results comparable to those above. Your *ping* results should show 0% packet loss and an equal number of packets transmitted and received. If some packets are being lost, the first thing you should check is the tightness and quality of the cable connections. Loose cables are frequently the cause of lost packets. The round-trip time factors are a function of the size and general load of your network, and not necessarily an indication of a problem with your Ethernet connection.

If your ping command is not successful, there are several things you should check. Perform these steps:

1. Try to ping the station by its IP address. For example, using our sample host *hancock*, enter the command

```
ping -r 128.70.3.56
```

2. Try to use the *ping* command on a different station on your local network, or use the broadcast address shown when you enter the command `ifconfig ec0`. You should see your system and answering duplicate responses from other systems.

3. Check the network configuration on your system with the `netstat -in` command. You should see information similar to this:

```
Name Mtu  Network  Address      Ipkts Ierrs Opkts Oerrs
ec0  1500  128.70.3 128.70.3.9  18    0    18    0
```

The `ec0` entry indicates your primary Ethernet connection. The `Ipkts` and `Opkts` fields indicate the number of inbound and outbound packets the network interface has processed. The `Ierrs` and `Oerrs` fields indicate the number of errors in input and output, respectively.

For the purposes of this troubleshooting session, though, check that the portion of the IP address shown under the `Network` heading match the IP address of the hostname that you attempted to `ping`. If the network addresses are not the same, the station is on a different network and the `ping` likely failed for that reason. Find a system on your immediate local network to which you can send the `ping` command.

4. Check the `/var/adm/SYSLOG` file for Ethernet error messages. Confirm that the designations for `netmask` and `broadcast`, as shown by `ifconfig ec0` match those of other systems on the same network or subnet.
5. Check to ensure that other stations are operating normally on the local network.
6. Check to ensure that the correct software (`eoel.sw.tcp`) package has been installed on your system.
7. Check the physical connections to the Ethernet cables for tightness and connection. If your Ethernet hardware is loose or disconnected, the `/var/adm/SYSLOG` file and your system console should both show messages such as

```
ec0: no carrier: check Ethernet cable
```
8. If all connections are tight and you still receive errors, replace the pieces of the Ethernet connection outside your system (the cable).
9. If you receive a message indicating that another host has the same IP address, find out which host has the same address, and determine which host has set their address incorrectly. (It is more likely that the same address was accidentally assigned to a second system, or that the new system being tested incorrectly set the address.)

Checking the Network Software Configuration

When the station comes up, verify the station's software configuration. Log into the station as *root*.

Enter the *chkconfig* command to check that your station's standard networking software is correctly configured:

```
/etc/chkconfig
```

You see information similar to this:

```
named      off
network    on
rwhod      off
timed      on
timeslave  off
gated      off
mouted     off
rtnetd     off
snmpd      on
routed     on
```

Note: Your output will vary from the output above depending upon installed software and configuration flag settings.

If you are familiar with the network-related daemons, you can customize your configuration flags to suit your network needs. If you are not familiar with the network-related daemons, set your network-related flags as shown above. In particular, make sure that the *network* variable is configured *on*.

About the Hosts Database

The *hosts* file is the hostname database. It contains information regarding known stations, from the perspective of the local station.

The */etc/hosts* database is an ASCII file that you can modify with any text editor. The file contains lines of text that specify a station's address, its "official" name, and any aliases. The "official" name should be the fully qualified domain name. The address and

name(s) are separated by blanks, tabs, or both. Comments begin with a pound sign (#) and continue to the end of the line.

Before you modify the hosts database, you should have a list of station names and valid Internet addresses of all stations in your network, as discussed in “IP Address Allocation” on page 36. If the network has routers (stations with multiple network interfaces), there must be a valid Internet address and name for each interface. See “Internet Protocol Addresses” on page 14 for a description of IP addresses.

This example assumes that you are not using NIS or BIND. If you are using NIS, refer to the *NIS Administration Guide* for more information. If you are using BIND, refer to Chapter 6, “BIND Name Server,” for more information.

An */etc/hosts* database is shown in this sample */etc/hosts* file:

```
# This is a comment
127.0.0.1    localhost
119.0.3.20  tuna.salad.com tuna  # tuna is an alias
119.0.3.21  chicken.salad.com  salad
119.0.3.22  walrus.salad.com  walrus
```

Each system must have a copy of */etc/hosts* that contains entries for **localhost** and all of its network interfaces. As shipped, the */etc/hosts* database contains two entries. The first entry is a name you can use to test the local network software:

```
127.0.0.1 localhost
```

When you reference **localhost**, the message is looped back internally; it is never transmitted across the network.

Caution: Many important programs depend on the **localhost** entry—*do NOT remove or modify it*. If the master copy of */etc/hosts* is not maintained on that system or if you are using BIND or NIS, make sure that the host database contains the **localhost** entry.

Note: The address 192.0.2.1 is a special address; networking will not be enabled if the system is given that address.

Modifying the Hosts Database

You need to edit two files to enable a system to access the network:

1. `/etc/hosts` file. You must add the system IP address in the `/etc/hosts` file as in the following example (see sample `/etc/hosts` above):

```
119.0.3.22 walrus.salad.com walrus
```
2. `/etc/sys_id` file. Your `/etc/sys_id` file entry *must* include the same system name as written in the `/etc/hosts`, whether as the official hostname or as an alias:

```
walrus
```

Here, host walrus can be expressed as either walrus or walrus.salad.com.

If you change the system's name in `/etc/sys_id`, update the entry in `/etc/hosts`; if not, the network software will not initialize properly. If the following message appears during station startup, then the `/etc/hosts` and `/etc/sys_id` files are inconsistent and must be fixed:

```
*** Can't find hostname's Internet address in /etc/hosts
```

If your system is a gateway, each network interface must be assigned an Internet address, each on a different network, and have an entry in `/etc/hosts`, as described in "Setting Up a Router" on page 47.

It is important that each station have a consistent version of the `host` database. The proper method for maintaining the consistency depends on the size of your network and whether the network is connected to the Internet. You can use the `rcp` or `rdist` programs by means of a `cron` job to ensure that the `hosts` files stay in sync.

Edit the `/etc/hosts` file and add the hostnames and Internet addresses for all stations on your network. Each station on the network must have all station names in the local `/etc/hosts` file. If you have a large `/etc/hosts` file, the easiest way to install it on a new system is to set up a minimal hosts file with entries for the new system and for another system that has an authoritative copy of the hosts file. This allows you to get the new system on the net and copy the more complete hosts file from the other system (using `rcp` or `ftp`). Another option is to transfer the hosts file on tape or disk.

Naming Your Station

Once you have determined your station's name, edit the `/etc/sys_id` file to give your station its new identity.

1. Remove the default station name (IRIS) and replace it with the new station name (setup1 for this example). You can use this command:

```
echo setup1 > /etc/sys_id
```

2. For the change to take affect, reboot your station with this command:

```
reboot
```

When your station comes back up, it should have the new station name as the login prompt.

Testing Your Network Connectivity

Two network management tools, *rup* and *ping*, provide quick information about network connectivity. *rup* indicates if there is a physical problem with the network, such as your station being unable to contact the other stations. Since *rup* uses broadcasts as a default, it does not go through routers. If your station can see the other stations on the network, use *ping* to test communication abilities. *ping* uses the Internet Control Message Protocol (ICMP), which requests an echo from the designated station. It lets you know if you can transmit and receive packets to and from specific stations.

1. Enter the *rup* command to determine if your station can contact the other stations on the network:

```
/usr/bin/rup
```

You should get output on each of the stations on your network. The other stations on your network must be up for your station to get a user-friendly response. If the other stations are powered on and attached to the network but not up in user mode, the information comes back in hexadecimal.

2. Enter the *ping* command to see if your station can communicate with the other stations on the network:

```
/usr/etc/ping station_name
```

Let the output run a few seconds, then use Ctrl+C to break it. Pay particular attention to the ping statistics. *ping* gives you the number of packets transmitted, number of packets received, percentage packet loss, and round trip time (minimum, maximum, and average). These are all good indicators as to the general condition of your network. Obviously, you want 0% packet loss and a fast round-trip time.

Setting Up a Router

A router is a station with multiple network connections that forwards packets between networks. This section provides the configuration procedure for a router with two interfaces and a router with more than two interfaces. A station can have multiple interfaces and not act as a router. The procedure for turning forwarding off on a station with multiple interfaces is also provided in this section. The specific subsections are

- “Configuring a Router With Two Interfaces” on page 47
- “Configuring a Router With More Than Two Interfaces” on page 48
- “Configuring Routing Behavior” on page 49
- “Turning On Multicast Routing” on page 50
- “About Multicast Packets Forwarding” on page 51
- “Setting Up Tunnels to Support Multicast Packets” on page 53

Configuring a Router With Two Interfaces

The `/etc/init.d/network` script is designed to automatically detect and configure a router with two interfaces *if* the default naming scheme for the interfaces is used. By default, the Internet addresses of the primary and secondary interfaces are derived from the `/etc/sys_id` file. The primary interface uses the name in the `sys_id` file. The secondary interface prefixes *gate-* to the name specified in the `sys_id` file.

To set up a router with two interfaces using the default naming scheme, follow this procedure:

1. Log in as **root**.

2. Assign valid Internet names and addresses to both interfaces in the `/etc/hosts` file. For example, the `/etc/hosts` file entries for the primary and secondary interfaces on the station `biway` might look like this:

```
198.70.75.2      biway.salad.com  biway
198.70.80.3      gate-biway.salad.com  gate-biway
```

3. Ensure that the router has the appropriate name in its `/etc/sys_id` file. Following this example, the `/etc/sys_id` file should look like this:

```
biway
```

4. Reconfigure the kernel and reboot the station to initialize your changes and interfaces. Some systems prompt you for permission, as in the following example. Others simply return a shell prompt. In either case, enter the `reboot` command when the kernel has been reconfigured:

```
/etc/autoconfig
Automatically reconfigure the operating system? (y/n)y
reboot
```

Note: If you do not want to use the standard router naming scheme, you must modify the `/etc/config/netif.options` file. “Modifying the Network Interface Configuration With the `/etc/config/netif.options` File” on page 56 details the procedure for changing an interface name.

Configuring a Router With More Than Two Interfaces

If the router contains more than two interfaces, you must modify the `/etc/config/netif.options` file in addition to the `/etc/hosts` and `/etc/sys_id` files. In the `netif.options` file, you must define the interface type (enp1, ipg0, and so on). By default, the names for the third and fourth interfaces are `gate2-$HOSTNAME` and `gate3-$HOSTNAME`, where `$HOSTNAME` is the value returned when you issue the `hostname` command. If you want to modify the interface names, see “Modifying the Network Interface Configuration With the `/etc/config/netif.options` File” on page 56 for the detailed procedure.

To set up a router with more than two interfaces and using the default naming scheme, follow this procedure:

1. Log in as **root**.
2. Assign valid Internet names and addresses to all interfaces in the `/etc/hosts` file. For example, the `/etc/hosts` file entries for the router *freeway*, with four interfaces, would look like this:

```
198.70.30.1    freeway
198.70.32.4    gate-freeway
198.70.41.5    gate2-freeway
198.70.59.6    gate3-freeway
```

3. Ensure that the router has the appropriate name in its `/etc/sys_id` file. Following this example, the `/etc/sys_id` file should look like this:

```
freeway
```

4. Modify the `netif.options` file to define your interface types. For this example, the third and fourth interfaces are FDDI (ipg*). Change the `if3name` and `if4name` variables from

```
if3name=
if4name=

to

if3name=ipg0
if4name=ipg1
```

5. Save your changes to the `netif.options` file.
6. Reconfigure the kernel and reboot the station to initialize your changes and interfaces. Some systems prompt you for permission, as in the following example. Others simply return a shell prompt. In either case, enter the `reboot` command when the kernel has been reconfigured:

```
/etc/autoconfig
Automatically reconfigure the operating system? (y/n)y
reboot
```

Configuring Routing Behavior

By default, when a station has two or more network interfaces, it automatically forwards (routes) packets between the two interfaces. If you do not want the station to serve as a router on its network, disable its route advertising.

1. Modify the `/etc/config/routed.options` file so the router will not supply routing information about general network routes (`-q`) or local interface routes (`-h`)

```
echo -qh > /etc/config/routed.options
```
2. Using the `/etc/init.d/network` script, turn the network off momentarily, then start it again

```
/etc/init.d/network stop  
/etc/init.d/network start
```
3. When the network comes up, verify that it is not forwarding packets with the `netstat` tool

```
netstat -s -p ip |grep forward
```

To change forwarding dynamics, create or edit the `/etc/config/routed.options` file to contain the desired options. Some of the behaviors that can be altered by means of the `/etc/config/routed.options` file are listed below:

- Suppress or force publicizing of routing information.
- Enable tracing of all received packets.
- Filter packets destined for specific networks.

See the `routed(1M)` online reference page for more details.

Turning On Multicast Routing

Multicast routing is a delivery method in which a message is sent to a designated group of receivers by the shortest delivery path (see Figure 3-1). SGI systems implement the Distance-Vector Multicast Routing Protocol (DVMRP, originally specified by RFC-1075) with the `mrouterd` process.

Note: As of IRIX 6.3, multicast routing can be used in combination with Network Information Service (NIS) to reduce the number of NIS servers needed throughout the network. For details, see the information on multicast-YP in the `portmap(1M)` and `nisserv(7P)` reference pages.

To turn on multicast routing:

1. Identify the routers on each network that need to support multicasting. Make sure that the routers you select are running IRIX Version 5.2 or later.
2. If you haven't already done so, install on each router the *eo.e.sw.ipgate* subsystem from your IRIX distribution source. Run *autoconfig* if necessary.
3. As **root**, enter the command
`chkconfig mouted on`
4. Restart the system with the *reboot* command.

About Multicast Packets Forwarding

Figure 3-1 shows an example network with three routers.




Figure 3-1 Network With Multicast Routers

- If people on networks A, C, and D are listening for packets, as shown in the figure, all four networks receive the multicast packets.
- If people on networks A and C are listening for packets, networks A, B, and C receive the multicast packets.
- If three or more people on network A are listening for packets, networks A and B receive the multicast packets. The multicast routing protocol prevents packets from being sent to “leaves” of the network, such as networks C and D; however, it doesn’t prevent packets from being sent to interior networks, such as network B.

Setting Up Tunnels to Support Multicast Packets

If your routers do not support multicast routing, you can support multicast packets by creating *tunnels* between the networks. Any SGI workstation running IRIX Version 5.2 or later can be configured as the endpoint of a tunnel. With tunneling, the multicast packets are encapsulated inside unicast packets, and then are sent to the other end of the tunnel. They are converted back into multicast packets when they are received.

Figure 3-2 shows an example of a setup with a tunnel between networks A and C.




Figure 3-2 A Tunnel Between Networks A and C

To create the tunnel, edit the file `/etc/mrouted.conf`. Step-by-step instructions follow:

1. Select the systems on each network that you will use for the sending and receiving end of the tunnel.

Choose a system that is running IRIX Version 5.2 or later, is fast, and is not used extensively. The audio and video data may be intermittent if the system you select is slow or overloaded.

2. If you haven't already done so, install the `oe.sw.ipgate` subsystem from your IRIX distribution source.

3. As `root`, edit the file `/etc/mrouted.conf` on the sending and receiving end of the tunnel. Note that these endpoints can be separated by many routers; you can use any machine on the network that is running IRIX Version 5.2 or later.

Add the following line for each network to which you want to establish a tunnel.

```
tunnel <local IP address> <remote IP address>
```

In the above example, the system on network D would have the following entry:

```
tunnel 128.70.58.1 128.65.170.2
```

The system on network A would have the following entry:

```
tunnel 128.65.170.2 128.70.58.1
```

4. You can specify other optional settings for the tunnel. For details, see the `mrouted(1M)` reference page.
5. Restart the system.

Note: One copy of the multicast packets is sent for each tunnel entry in `mrouted.conf`. This results in extra network traffic. For example, suppose you have a workstation with one network interface and you set up tunnels to workstations on three other networks. The packets are replicated three times.

Updating `/etc/rpc` for NIS Users

The IRIX copy of the `/etc/rpc` file contains additions that are essential for using certain multicast utilities if you're running the Network Information Service (NIS). If the NIS master is not running IRIX Version 5.2 or later, or is not a SGI workstation, verify that the `/etc/rpc` file on the NIS master includes these additions:

```
sgi_iphone 391010  
sgi_videod 391011
```

Subnetting a Network

Implementing the subnet address scheme is a very simple procedure. Subnetting consists of the following:

1. “Setting the Netmask” on page 55
2. “Rebooting the Station” on page 56

Note: If you have not done so already, read Chapter 2, “Planning a Network.” It contains information about partitioning Internet addresses for subnetting.

Setting the Netmask

The **netmask** option of the *ifconfig* command is used to interpret and define the network portion (subnets included) of the Internet address. A *network mask* defines the portion of the Internet address that is to be considered the network part. This mask normally contains the bits corresponding to the standard network part as well as the portion of the host part that has been assigned to subnetworks. If no network mask is specified when the address is set, it will be set according to the class of the network.

To configure a station’s primary interface to recognize a subnetted Class B address that is using 8 bits of the host ID for defining the subnets, create or modify the */etc/config/ifconfig-1.options* file and insert the following line:

```
netmask 0xfffff00
```

This netmask value indicates that for the primary interface, the upper 24 bits of the Internet address represent the network portion of the address, and the remaining 8 bits represent the host ID. The nonsubnetted netmask for a Class B address (in hexadecimal) would be 0xffff0000. The netmask value can be set using hexadecimal, dot-notation Internet address, or pseudo-network name formats. This entry must be made on all stations on the subnet.

Note: For stations with multiple interfaces, the network mask should be set for each interface in the appropriate *ifconfig-*.options* file.

Although variable subnet masks, that is, different masks on different interfaces, are supported as of IRIX 6.2, they are difficult to set up and maintain. It is therefore strongly recommended that you avoid using them unless there is no other way of accomplishing your site requirements.

Rebooting the Station

When the **netmask** value is set, the station must be reconfigured and rebooted to incorporate the new network address into the stations routing tables. Always reboot routers before regular stations, because they provide routing information to other stations and networks.

Reconfigure the kernel and reboot the station to initialize your changes and interfaces. Some systems prompt for permission before reconfiguring the kernel, while others simply return a shell prompt. In either case, enter the *reboot* command explicitly:

```
/etc/autoconfig
Automatically reconfigure the operating system? (y/n)y
reboot
```

Modifying the Network Interface Configuration With the */etc/config/netif.options* File

You do not always need to modify (configure) a station's network interface; in most situations, the default configuration suits the site's needs. Modifying the network interface configuration requires that you modify the */etc/config/netif.options* file. You modify this file if

- the station has more than two interfaces
- you don't like or use the default naming conventions
- the default order is incorrect, or you prefer a different order.

You can identify the station's installed network interfaces using the `hinv` command with the `network` option, as described on "Controller Interface Names" on page 4. See also the `hinv(1M)` reference page. For a summary of configuring multiple network interfaces see "Configuring Multiple Network Interfaces" on page 67.

There are two configurable variables in the `netif.options` file: *interface name* and *interface address*. The interface name variable designates the order (first, second, third, or fourth) and type of interface involved. The interface address variable assigns a valid Internet address to each interface. There must be a valid Internet address for each interface in the `/etc/hosts` file.

These subsections tell you about the following tasks

- "Modifying the Interface Name in the `/etc/config/netif.options` File" on page 58
- "Modifying the Interface Address in the `/etc/config/netif.options` File" on page 59

Table 3-1 summarizes the interface name and interface address variables.

Table 3-1 Variables for the `netif.options` File

Variable Name	Variable	Examples
Interface Name	<code>ifxname=</code> where: <code>x = 1, 2, 3, or 4</code> <code>name=ec0, et0, enp0, enp1, fxp1, ipg0, ipg1, etc.</code>	<code>if1name=enp0</code> <code>if2name=ipg0</code> <code>if3name=enp1</code> <code>if4name=enp2</code>
Interface Address	<code>ifxaddress=</code> where: <code>x = 1, 2, 3, or 4</code> <code>address=\$HOSTNAME, station name, or Internet address</code>	<code>if1address=\$HOSTNAME</code> <code>if2address=fddi-\$HOSTNAME</code> <code>if3address=gate-goofy</code> <code>if4address=128.70.28.2</code>

You can modify either or both variables. Instructions for modifying both variables are provided below.

Note: `$HOSTNAME` will be set to the value in `/etc/sys_id`.

Modifying the Interface Name in the `/etc/config/netif.options` File

When a station has more than two network interfaces, you must modify the name entries in the `/etc/config/netif.options` file to assign the interface order. Default order is assigned only to the first two interfaces. Additionally, if you want to change the order (first, second, and so on) or interface type assigned to a network interface, you must modify the `/etc/config/netif.options` file. This example makes the first FDDI interface secondary.

1. Using the `netstat` command, verify the network interface's name:

```
/usr/etc/netstat -ina
```
2. Using `vi` or any editor, open the `netif.options` file for editing:

```
vi /etc/config/netif.options
```
3. Locate and modify the appropriate interface name variable. For this example, search for the secondary interface name variable (`if2name`) and change it from the default configuration to a configuration that supports the first FDDI interface as secondary.

Change from this

```
: if2name =
```

to this

```
if2name=ipg0
```

Caution: Note that all *default* variables (primary and secondary) start with a leading colon (:). You must remove it *and* enter the interface type to change the default interface name.

4. Save and exit the file.

If you have no other changes, `autoconfigure` and reboot the station. Otherwise, repeat this procedure for each interface name change.

Note: When you alter the order of one network interface, the other interfaces in the station remain in their default relative ordering. For example, on a three-interface station (a=first, b=second, and c=third), if you were to make the default second the first (b=first), the remaining interfaces would be configured a=second and c=third.

Modifying the Interface Address in the `/etc/config/netif.options` File

To change the default interface address, modify the `/etc/config/netif.options` file. All interface names require valid Internet addresses as found in the `/etc/hosts` file. The `$HOSTNAME` variable pulls the station name from the `/etc/sys_id` file. This example changes the second, third, and fourth interface addresses as follows:

- second interface address: `fddi-$HOSTNAME`
- third interface address: `gate-goofy`
- fourth interface address: `128.70.28.26`

Follow these instructions to modify your network interface address according to the above example:

1. Verify or assign a valid entry for each interface in the `/etc/hosts` file. Write down the name and address for each interface.
2. Using `vi` or any editor, open the `netif.options` file for editing:

```
vi /etc/config/netif.options
```
3. Locate and modify the appropriate interface address variable. For this example, we want to modify the second, third, and fourth interface address variables. Find and modify each variable as follows.

Change from this

```
: if2addr=gate-$HOSTNAME
if3addr=gate2-$HOSTNAME
if4addr=gate3-$HOSTNAME
```

to this

```
if2addr=fddi-$HOSTNAME
if3addr=gate-goofy
if4addr=128.70.28.26
```

Caution: Note that all *default* variables (primary and secondary) start with a leading colon (:). You must remove it *and* enter the interface address to change the default interface address.

4. Save and exit the file.

Repeat this procedure for each interface address change. If you have no other changes, reconfigure and reboot the station.

Assigning IP Aliases

You can assign multiple IP addresses to a single network interface using the IP aliasing feature of the `ifconfig` command (see `ifconfig(1M)`). This feature is useful if you find it necessary to use the same network interface for connections to several networks; it is also useful for maintaining an obsolete address for a host until users learn its new address. IP aliasing is supported on all network interface types.

The following subsections will help you understand and implement IP aliasing:

- “Manually Implementing IP Aliasing” on page 60
- “Guidelines for Assigning IP Aliases” on page 61
- “Creating an IP Aliases File” on page 61
- “IP Aliasing and Routing” on page 62

Manually Implementing IP Aliasing

When IP aliasing is implemented on a host, the address that the `/etc/host` file assigns to the hostname in `/etc/sys_id` is considered the *primary address* for the interface, and alternative names assigned by the `ifconfig` command are considered *IP aliases*. When you display network information (see “Collecting Network Statistics With `netstat`” in Chapter 4), the primary address of an interface is listed first, followed by its aliases.

To assign an IP alias, use this form of the `ifconfig` command:

```
# ifconfig interface alias address [netmask mask_num] [broadcast address]
```

Note: Be sure to read the information in “Guidelines for Assigning IP Aliases” (below) before assigning IP aliases.

You can specify a hostname from `/etc/hosts` instead of an IP address as the value of *alias*. For example, either command below assigns an IP alias to the `ec0` interface. The interface in this example uses a netmask and broadcast address, but these fields are optional:


```
# ifconfig ec0 alias 128.70.40.93 netmask 0xffffffff broadcast 128.70.40.255
# ifconfig ec0 alias plum netmask 0xffffffff broadcast 128.70.40.255
```

Guidelines for Assigning IP Aliases

In theory, the number of IP aliases assigned to a given network interface is limitless; however, the recommended maximum number is 1000. In addition, the primary address of an interface and all its IP aliases should share a common subnet address. For example, if the primary address of an interface on a class B network is 128.70.12.23, its IP aliases should be addresses such as 128.70.12.19, 128.70.12.53, and so on.

Creating an IP Aliases File

You can save IP alias settings in the `/etc/config/ipaliases.options` file so they are automatically assigned during the network initialization process. The `ipaliases.options` file is read by the network script if the `/etc/config/ipaliases` file contains the value `on`. Use the procedure below to create an `ipaliases.options` file and assign IP aliases automatically during network initialization:

1. Edit `ipaliases.options` using your favorite editor:

```
# vi /etc/config/ipaliases.options
```

Your entries should look similar to the example below. In this example, broadcasting is performed on the Ethernet but not the FDDI network:

```
ec0 128.64.56.51 netmask 0xffffffff broadcast 128.64.56.255
ec0 128.64.56.12 netmask 0xffffffff broadcast 128.64.56.255
ipg0 190.111.79.16 netmask 0xffffffff
ipg0 190.111.79.10 netmask 0xffffffff
```

2. Create an `ipaliases` file and set IP aliasing on in the file:

```
# chkconfig -f ipaliases on
```

3. Stop and restart the network to put your changes into effect:

```
/etc/init.d/network stop
/etc/init.d/network start
```

IP Aliasing and Routing

Each IP alias is maintained in the IRIX routing table as a host route between the primary address and the alias. By default, IRIX assumes that a host containing multiple network interfaces is a router. These *multi-homed hosts* perform packet forwarding and advertise routing information on the networks to which they are attached. However, if you implement IP aliasing on a host with a single network interface (to connect a local network and the Internet, for example), route advertising is not enabled. In this case, it is necessary to force route advertising to make the network location of the local hosts known.

To force route advertising, enter the `-s` option in the `routed.options` file on the IP-aliased host.

Changing Network Parameters in the `ifconfig-#.options` File

Network parameters are those settings that determine how an interface processes or supplies certain network information. Modifying network parameters requires that you create and modify the appropriate `/etc/config/ifconfig-#.options` file (where # can be 1, 2, 3, or 4, based on the interface name).

The default parameter settings are known to function well and suit most site needs. Changing the settings can cause a network to become dysfunctional. It is recommended that these parameters be changed only by experienced or trained network managers.

There are four nondefault network parameters that can be set in the `ifconfig-#.options` file:

- netmask
- broadcast address
- Address Resolution Protocol (ARP)
- route metric

These steps show how to set the nondefault network parameters:

1. Using the `netstat` command, determine the order assigned to the network interface at configuration time:

```
/usr/etc/netstat -i
```

- Using `vi` or any editor, open or create the file `/etc/config/ifconfig-#.options`, where the pound sign (#) represents the network interface's name. For example, to configure a primary interface, open or create the file `/etc/config/ifconfig-1.options`.
- To change the network mask value for a network interface, enter a line with the word `netmask` followed by a space and either the 32-bit hexadecimal value, Internet address dot-notation, or the pseudo-network name:

```
netmask 0xffffffff00
```

See "Turning On Multicast Routing" on page 50 for more details regarding netmasks.

- To change the broadcast address for a network interface, enter a line with the word `broadcast` followed by a space and the dotted decimal IP broadcast address:

```
broadcast 189.170.6.0
```

To enable or disable the Address Resolution Protocol (ARP), enter a line with `arp` (to enable) or `-arp` (to disable):

```
arp
```

The ARP tables are used by some of the network management tools (`netstat`, `ifconfig`, and so on) and provide the administrator with invaluable network information.

- To change the routing metric count for a network interface, enter a line with the word `metric` followed by a space and the count:

```
metric 7
```

The default metric count, also called *hops*, is 0. The `routed` daemon monitors the number of hops required to route data from one network to another. If you want to reduce network traffic on a particular route, increase the metric count on the appropriate router.

The interface configuration file for the secondary interface might look like the following:

```
cat /etc/config/ifconfig-2.options
```

```
netmask 255.255.255.0  
broadcast 129.78.50.0  
-arp  
metric 4
```

The interface configuration file above indicates that the Class B network is subnetted using 8 bits of the host ID for the subnet. It uses 0 as the broadcast address instead of the

default 1. ARP has been disabled and the metric (hop) count has been set to 4 to discourage router traffic.

Configuring `/etc/gateways` Files for Networks That Do Not Support Broadcast or Multicast

Note: A *gateway* is a host or router on which the *routed* daemon runs with the `-s` option to supply routing information. A *client*, on the other hand, does not supply routing information. Therefore, it runs *routed* with the `-q` option. (See the *routed(1M)* reference page for more information.)

The `/etc/gateways` file is a list of distant gateways (that is, hosts and routers that supply routing information) that is used by networks that do not support broadcast or multicast, such as ATM or HIPPI.

This file tells a host either of two things:

- Which gateways have hosts on their subnets
- Which subnets/gateways are reachable through a directly-connected gateway

Hosts then use this file to determine where to send routing information protocol (RIP) requests and, if the hosts also supply routing information, where RIP responses should be sent.

`/etc/gateways` File Format

Entries in the `/etc/gateways` file should include all gateways reachable on the ATM or HIPPI network in the following formats:

The `/etc/gateways` file includes a series of lines, each in one of the following two formats.

Format 1:

```
host Nname [mask value] gateway Gname metric value [passive|active|external]
```

Format 2:

```
host Hname gateway Gname metric value [passive|active|external]
```

These are described as follows:

<i>Nname, Hname</i>	The name of the destination network or host.
<i>mask value</i>	Optional. Value between 1 and 32 to indicate the netmask associated with <i>Nname</i> .
<i>Gname</i>	The address of the gateway to which RIP sends requests and, should the host provide routing information, responses.
<i>metric value</i>	The hop count to the destination host or network.
<i>passive</i>	The gateway is not expected to exchange RIP packets. Routes through passive gateways are installed in the kernel's routing tables at startup and are not included in routing updates.
<i>active</i>	Indicates if the gateway is willing to exchange RIP packets.
<i>external</i>	Indicates that the gateway is passive, but placed in the kernel routing table. Entries marked as external indicate that another routing process will install such a route, if necessary.

Lines that do not start with `net` or `host` must consist of one or more of the following parameter settings, separated by commas or blanks:

<code>if=ifname</code>	This parameter indicates that the other parameters on the line apply to the interface name <code>ifname</code> .
<code>subnet=nname</code> <code>[/mask value][,metric]</code>	This parameter advertises a route to the network <i>nname</i> with the <i>mask value</i> and the supplied <i>metric value</i> (default is 1). The network number must specify a full, 32-bit value, as in <code>192.0.2.0</code> instead of <code>192.0.2</code> . If you choose to add an alias to the loopback interface, the alias must be reachable.
<code>passive</code>	This parameter marks the interface to not be advertised in updates sent via other interfaces, and turns off all RIP and router discovery through the interface. This is very useful for failsafe interfaces or for interfaces for which you do not want RIP traffic to route through.
<code>no_ripv1_in</code>	This parameter causes RIPv1-received responses to be ignored.
<code>ripv2_out</code>	This parameter generates RIPv2 output, and causes RIPv2 advertisements to be multicast, when possible.
<code>ripv2</code>	This parameter is equivalent to <code>no_ripv1_in</code> and <code>ripv2_out</code> .

```
rdisc_interval=N
```

This parameter sets the nominal interval with which router discovery advertisements are transmitted to N seconds and their lifetime to $3*N$. This should be set to 45.

Also, rather than running *routed* with the **-P parm** options, you can instead add the parameter line *parm* to the */etc/gateways* file.

Examples of */etc/gateways* Files

The */etc/gateways* file in Example 3-1 belongs to a host that has two HIPPI interfaces, one on the `192.168.0.0/27` subnet and the other on subnet `192.168.0.96/27`. This host functions as a client. It sends RIP requests to one gateway on the `192.168.0.0/27` subnet, and to two gateways on the `192.168.0.96/27` subnet. It should neither generate nor process RIP packets on its `et0` interface.

Example 3-1 */etc/gateways* File for Host With Two HIPPI Interfaces

```
#if broadcast is not supported
#-- 192.168.0.0/27 subnet
host 192.168.0.1 gateway 192.168.0.1 metric 1 active i
#-- 192.168.0.96/27 subnet
host 192.168.0.97 gateway 192.168.0.97 metric 1 active
host 192.168.0.99 gateway 192.168.0.99 metric 1 active

ripv2_out
rdisc_interval=45
if=et0 passive
if=et1 passive
```

The */etc/gateways* file in Example 3-2 belongs to a host that functions as a gateway. It has two HIPPI interfaces, one on subnet `192.168.0.96/27` and one on subnet `192.168.0.128/27`. There is one other gateway with which this host should exchange routing updates on the `192.168.0.96/27` subnet, and one on the `192.168.0.128/27` subnet.

This host solely and permanently owns an IP address in the `192.168.0.202` subnet; this IP address is actually an alias that we want advertised as long as there is at least one active interface left, thus the `subnet=192.168.0.202/31,1` entry. It should neither generate nor process RIP packets on its `et0` interface.

Example 3-2 `/etc/gateways` File for Host Functioning as a Gateway

```
#if broadcast is not supported
#-- 192.168.0.96/27 subnet
host 192.168.0.97 gateway 192.168.0.97 metric 1 active
host 192.168.0.99 gateway 192.168.0.99 metric 0 active
#-- 192.168.0.128/27 subnet
host 192.168.0.129 gateway 192.168.0.129 metric 1 active
host 192.168.0.131 gateway 192.168.0.131 metric 0 active

subnet=192.168.0.202/31,1

ripv2_out
rdisc_interval=45
if=ef0 passive
```

For further information, see the `routed(1M)` reference page.

Configuring Multiple Network Interfaces

Configuring a station with multiple network interfaces is much like configuring a router, except that forwarding is not enabled. With several network controllers (physical boards or chips), the system must configure a unique interface name for each. Multiple interfaces usually connect systems to more than one subnet and require an address for each connected interface.

The default network configuration file is `/etc/init.d/network`, which initializes and tests the network interface names and addresses. The site-specific network configuration files dealing with the daemons and their options are `/etc/config/netif.options` and `/etc/config/ifconfig-<n>.options` (where `<n>` is 1, 2, 3, and so on, based on the interface variable name). The file `/etc/config/netif.options` overrides the defaults specified in the file `/etc/init.d/network`.

By default, the network configuration script only configures two network interfaces, regardless of how many network interfaces exist on a system. To increase the number of configured network interfaces you must follow these steps:

1. Identify all the unique network interfaces—“Controller Interface Names” on page 4.
2. Modify the network interface configuration in the `/etc/config/netif.options` file—“Modifying the Network Interface Configuration With the `/etc/config/netif.options` File” on page 56.

3. Assign and implement IP aliasing—“Assigning IP Aliases” on page 60.
4. Change network parameters in the *ifconfig-#.options* file—“Changing Network Parameters in the *ifconfig-#.options* File” on page 62.
5. Reboot the system.

Dynamic Host Configuration With Proclaim

This section describes the IRIX *proclaim* facility, which is based on the dynamic host configuration protocol (DHCP) described in RFC 2131. Proclaim and DHCP allow you to allocate hostnames and network addresses automatically.

DHCP permits site administrators to set up one or more server systems that dynamically distribute network IP addresses and site configuration parameters to new or requesting client systems. In this way, a site with only a few available addresses can serve a large number of hosts that connect to the network only occasionally, or a large site can manage the permanent assignment of addresses with a minimum of administrative attention.

The *proclaim* facility consists of a daemon that runs on a master server, optional relay daemons for servers on subnets, GUI configuration programs, and a client application that communicates with the DHCP servers. All these programs are bundled with the IRIX operating system.

If you need to refresh your knowledge of the basic concepts of networking, IP addresses, and netmasks, please read “Internet Protocol Addresses” on page 14, and “Setting the Netmask” on page 55 before continuing with this section.

These subsections provide information about configuring *proclaim* servers and clients:

- “DHCP Server Basic Operation” on page 69
- “Configuring the DHCP Server” on page 69
- “Configuring the DHCP Relay Agent” on page 83
- “Setting up the DHCP Proclaim Client” on page 85
- “Limitations of DHCP” on page 90

DHCP Server Basic Operation

The DHCP server is started by *inetd*(1M), as configured in the *inetd.conf* file. The basic operation of the DHCP protocol is a series of packet exchanges as follows:

1. The client machine broadcasts a DHCPDISCOVER packet to the DHCP server's UDP port, using a UDP broadcast equivalent. The request packet includes the requester's network hardware address, and a list of desired configuration options.
2. All the DHCP servers on the same network as the client machine receive the client's request as do the *bootp* relay agents, which forward the request to other subnets.
3. The server looks up the requester in its configuration files to determine if it is configured to serve clients from the requesting subnet. The server now looks in the *propel* database or the */etc/ethers* file (see *ethers*(4)) for an entry with the client's network hardware address. If an entry is found, the server checks the hostname of that entry against the *propel* database or the */etc/hosts* file (see *hosts*(4)) to complete the network hardware address to Internet address mapping.

If no entry is found, the server generates an appropriate IP address and hostname using its configuration files; see *dhcp_bootp*(1M).

4. The server now returns a response in the form of a DHCPOFFER packet, which also includes other configuration information requested by the client.
5. The client accepts the first valid offer from a server and broadcasts back a DHCPREQUEST message, which includes all the configuration options from the DHCPDISCOVER message plus the address of the selected server.
6. The selected server makes a permanent mapping for the client's hostname, IP address, and hardware address, and responds with a DHCPACK message.
7. On receiving the DHCPACK message from the server, the client writes the configuration information into the appropriate configuration files and then executes the networking initialization scripts.

Configuring the DHCP Server

The *dhcp_bootp* program is a server that communicates with DHCP and *proclaim* clients to provide host configuration information, including an IP address at minimum. If your site uses DNS to maintain the hosts map, you need some external mechanism (such as a script) to update the DNS map from the DHCP server. If your site uses NIS to

maintain the hosts and ethers maps, *dhcp_bootp* must be run on the same machine as the NIS master server.

The DHCP server generally uses configuration parameters based upon the subnet number of the originating client request. The configuration files are all placed in the directory */var/dhcp/config* and are named in the form *config.<netnumber>*. For example, the configuration file for clients on the 128.78.61 network should be named *config.128.78.61.0*. If the subnet configuration file does not exist, then the *config.Default* file applies instead. If the default configuration file is missing or unreadable, clients are configured to match the DHCP server itself.

To configure the DHCP server, you can do one of the following:

- Run the *ProclaimServerMgr* graphical interface, as described by the *ProclaimServerMgr(1M)* reference page.
- Edit the */etc/config/dhcp_bootp.options* file and use the *dhcp_bootp* command to control the execution of the DHCP server

ProclaimServerMgr

ProclaimServerMgr is a graphical front end that displays the status and configuration of the *proclaim* server running on the local host. It displays the set of IP addresses and lease times that are currently assigned. *Proclaim* is an implementation of DHCP (dynamic host configuration protocol). Use the following command to invoke *ProclaimServerMgr*.

```
/usr/sbin/ProclaimServerMgr
```

Figure 3-3 shows the graphical interface for the *proclaim* server configuration.

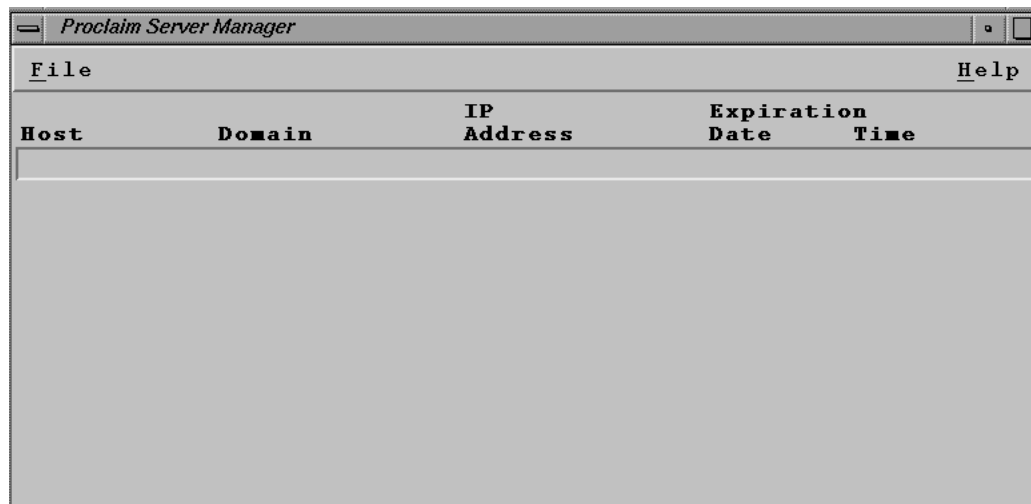


Figure 3-3 ProclaimServerMgr Graphical User Interface

The *File* button on the main window opens a pull-down menu for editing options and configuration files as show in Figure 3-4.

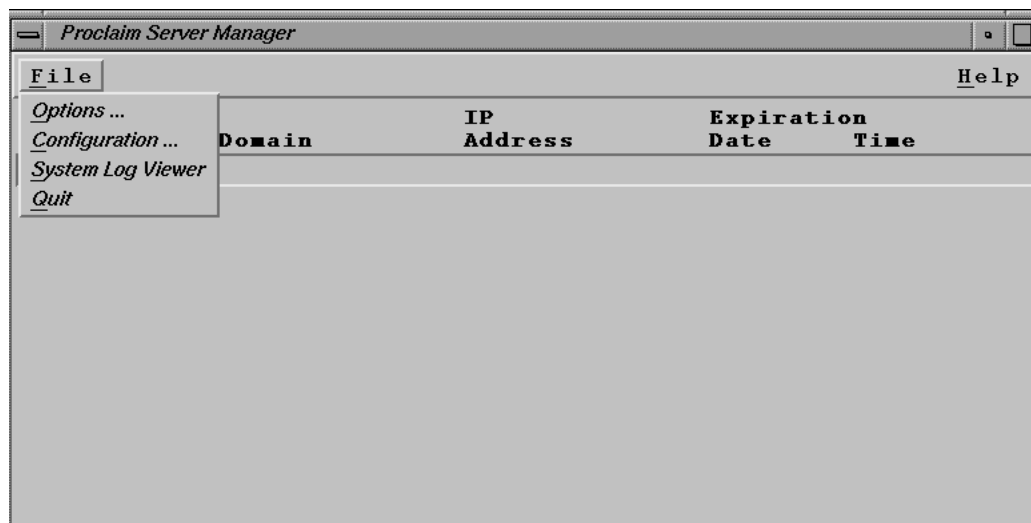


Figure 3-4 ProclaimServerMgr File Button Pulldown Menu

The *Options* menu choice allows you to edit proclaim server options defined in the `/etc/config/proclaim_server.options` file.

The *Configuration* menu choice allows you to add, delete, and edit configuration parameters. Configurations are saved in files and define the IP addresses, lease times, and other options returned to clients. See "*DHCP Server Configuration Parameters*" on page 78, and the `dhcp_bootp(1M)` for details about configuration parameters.

To view messages logged by the server in the system log file (`/var/adm/SYSLOG`), launch `sysmon(1M)` from the main window by choosing the System Log Viewer menu item.

You can disable or enable DHCP by clicking the Proclaim Server **on/off** button in the main window.

Using the `dhcp_bootp` Command to Configure the DHCP Server

You can use the `dhcp_bootp` command to control the execution of the DHCP server.

You can specify the additional DHCP-related options in the `/etc/config/dhcp_bootp.options` file. The command is as follows:

```
/usr/etc/dhcp_bootp [ -fd ] [ -o dhcp_options_file ]
```

The `dhcp_bootp(1M)` command accepts the following options:

- f This enables cross network forwarding if the client is requesting for service from a dhcp server which is reachable via this server and on a different wire.
- d This enables the `dhcp_server` into a debug mode in which it logs messages to `/var/adm/SYSLOG`.
- o The path of the `dhcp_options_file` can be mentioned using this option. The default location is `/etc/config/dhcp_bootp.options`.

Controlling the Execution of the `dhcp_bootp`

You can control the execution of the `dhcp_bootp` program by following these steps:

1. Use vi or your some other text editor to specify the additional DHCP-related options in the `/etc/config/dhcp_bootp.options` file as follows:
 - y
If you are using NIS to propagate the hosts and ethers map, this system is required to be the NIS master server for proclaim to work correctly.
 - w *hosts_map*
Specifies the optional location of the hosts map. The default is `/etc/hosts`.
 - e *ethers_map*
Specifies the optional location of the ethers map. The default is `/etc/ethers`.
 - u *sysname*
Specifies the name for an optional *sysname* file. The default is `/unix`.
 - c *proclaim_config_dir*
Specifies the optional proclaim server configuration directory. The default directory is `/var/dhcp/config`
 - t *seconds*
A timeout period for the *dhcp_bootp* server can be set else the default period will be used. If a timeout period of 0 seconds is used then the *dhcp_bootp* server will become a persistent daemon.
 - x
This option is used in conjunction with a Domain Name Server that provides the IP address to hostname mapping for the initial assignment of IP addresses to hosts by the DHCP server. The option has been found to be particularly useful in an Internet Service Provider type of setup where the IP address to hostname mapping may be required to remain fixed. Once an IP address is assigned it is written to the *etherToIP.pag* file. When this option is in effect, *etherToIP.pag* is the only file consulted to determine whether a particular IP address or hostname is already assigned.
 - n
Specifies that the server should return a netmask with its reply regardless of whether a client requested it or not.

- `-m options` Specifies additional options to be sent by the server to clients irrespective of whether the client requests them or not. The additional options are specified as a comma or colon separated list with no spaces.
- `-r script_file_path` Specifies an optional script file to execute when the server changes the etherToIP database. The complete path to the script must be specified as an a argument. The server calls the script with several command-line arguments: `-c client_identifier -m physical_address -i ip_address -h hostname -l lease -o operation -t dhcp_message_type`. The script is executed in the background. The value of the operation argument is 0 if it was a record creation, 1 if it was a deletion, and 2 if the lease was updated. The `dhcp_message_type` is the same as defined message types in the DHCP literature. If a message is sent to the client that message type is passed as an argument, otherwise the message type received from the client is passed. This script may be used to perform other functions such as dynamic DNS updates.
- `-W` Specifies that the the hosts map file or `/etc/hosts` file should not be updated when leases are added or removed.
- `-E` Specifies that the ethers map file or `/etc/ethers` file should not be updated when leases are added or removed.
- `-pn` Specifies that a ping be sent before an address is leased to check if a host is using the IP address. If a reply is received then a message is logged in the SYSLOG indicating that "Address <addr> appears to be STOLEN" and a record is added into the database with a lease time of -2 for this IP address. The DHCP administrator should investigate the occurrence and remove the entry in the database with the `dhcpdb` command once it has been resolved. The DHCP server tries to give out another address to the requesting client. The server attempts to recycle stolen addresses when it runs out of addresses and at a frequency specified by the lease time in the configuration files.

`-pt ping_timeout`

Specifies the timeout for the ping. If a reply to a ping is not received in *ping_timeout* seconds the address is considered to be not taken. The default value for the timeout is 1 second. Note that the server processes requests from other DHCP clients while waiting for the timeout to expire. The server does not block waiting for a reply to a ping request. If blocking pings are required for some reason use the `-pb` option instead of the `-pn` option.

`-pl number_of_outstanding_pings`

This is the maximum number of pings outstanding. Once this value is reached the server automatically stops sending ping to check if an IP address is taken. The server begins to send ping request once the number of outstanding ping requests falls below his value. It is set to 16 by default.

`-pd`

This option is available for backward compatibility with earlier versions. In prior versions (if `-x` was not specified) while assigning a new address the server would make a *gethostbyaddr(3N)* call to check if it could assign an IP address to a client for which it had no prior binding. If a binding was found in DNS maps then the particular address was not assigned. In the new version of the server to obtain this older behavior it is necessary to specify the `-pd` option. If this is not specified the server will not perform a *gethostbyaddr(3N)* call to check if an IP address is available.

The following options (`-dx,-l`) can be set only with the Enhanced DHCP server product.

`-dn`

Specifies that dynamic DNS updates are on. The behavior of DNS updates are modified by the following (`-da-db-dc-ds-dt`) additional options. Note that if the `-x` DHCP server option is specified then the `-dn` option is silently ignored. With the `-x` option DNS maps are pre-populated so it does not make any sense to update DNS dynamically.

`-da`

Ordinarily the server will perform dynamic DNS updates (if the `-dn` option is specified) only if the Dynamic DNS option is sent by

the client (option 81 as per the draft). If this (`-da`) option is used the server will update DNS regardless of whether option was sent by the client or not.

`-db`

If this is specified the DNS updates are performed before an ACKnowledgement is sent to the client. By default the dynamic DNS updates are performed after the ACK is sent to the client.

`-dc`

If this is specified the DNS A RR is also updated by the server. Otherwise whether the A RR is updated or not depends on the flags value sent in the dynamic DNS update option by the client.

`-ds`

Specifies that the security feature should not be used. By default secure updates are performed using TSIG for those networks for which a key is specified.

`-df filename`

Specifies the file that contains the (TSIG) keys for the various networks served by the DHCP server. The default file name is `/var/dhcp/config/dhcp_ddns.conf`. A key must be specified for each network number for which secured updates are required. Each line in this file specifies the network number (netnumber) followed by the keyname, algorithm, and the base 64 encoded secret key. Fields are separated with spaces or tabs. The keyname and algorithm-name are the same as that used in the BIND 8.2 configuration file.

`-dt time_to_live`

Specifies the time to live (ttl in seconds) for records that are added to the DNS. The ttl specified by the DHCP server in the DNS update is the minimum of this value and the lease duration.

`-l /var/dhcp/config/dhcp_ldap.conf`

This option specifies that the ldap related options should be picked from the `/var/dhcp/ldap/dhcp_ldap.conf` file. See the file `/var/dhcp/ldap/dhcp_ldap.conf` to see various options and what may need to be specified in order to use the DHCP server with an LDAP server to store configuration and leases.

2. Create one or more network configuration files based on the network numbers this server is responsible for. A default configuration file is provided, which is `/var/dhcp/config/config.Default`. The new file(s) should be named `config.<netnumber>` in the same directory. For example, if this server is serving clients on the 192.26.61 network, the configuration file should be named `config.192.26.61.0`. See the list of configuration parameters below.
3. The DHCP server also provides MAC address based filtering. The filters are specified in the `/var/dhcp/config/dhcpmtab` file. MAC addresses can be specified as either complete addresses or as a regular expression to specify a set of addresses. The regular expression syntax is described in the man pages for `regcmp(3G)`.

If a DHCP packet from a client with a matching MAC address is received then the corresponding actions are taken. Regular expression matching is attempted only if a complete match is not found. The length and hardware type must also match for non-regular expression entries. The length may not match for entries that have regular expressions. Once a match is found the remaining entries are not searched. An example of the configuration file is as follows:

```
# dhcpmtab - mac address based filtering for DHCP
#len htype      mac                action
6    1    01:02:03:8a:8b:8c    K
0    1    01:02:*                  L|K
```

The actions can be specified as a combination of L and K. Each of the actions L=log request, and K=drop request, is acted upon in that order. The system log is updated if the L action is requested. If the K action is specified and an L is not specified then the packet is silently dropped.

4. The DHCP server can be configured to send specific user defined options based on the value of the vendor class option sent by the DHCP client. To define new options for this purpose use the `/var/dhcp/config/dhcp_option_types` file. The format of this file is described as comments within this file. Each line consists of a user selected name, an option number, and the type of the option. The value to be returned by the DHCP server can be specified in the `/var/dhcp/config/vendor_options` file. See the file for the format. The vendor class and the network number to which option values apply can be specified within the file.
5. Enable the proclaim server by giving the command:


```
chkconfig proclaim_server on
```

6. Replace the "bootp" entry in `/usr/etc/inetd.conf` with this line:

```
bootp dgram udp wait root /usr/etc/dhcp_bootp dhcp_bootp -o \  
/etc/config/dhcp_bootp.options
```
7. Issue this command, to restart inetd with `/etc/killall -HUP inetd`:

```
/etc/killall bootp
```

DHCP Server Configuration Parameters

These parameters can be supplied in the configuration file for each network. All of the address fields can either be in standard IP address dot notation, or as a hex number prefixed with a 0x. Most of the fields can be left blank to render them non-applicable.

`pro_address_counter`

This integer field specifies the host number for the next IP address. The next address is constructed using the counter and checked through the range of the assignable addresses. The first available address in the range is assigned.

`pro_host_pfx_counter`

This integer field specifies the starting number appended to the `pro_host_prefix` to generate a new hostname. This counter is incremented and a new hostname generated until a unique unused name is found.

`pro_netmask`

This field takes a netmask in address form. For more information on netmasks, see "Setting the Netmask" on page 55. This field specifies the subnetmask used by the client systems.

`pro_lease`

This unsigned integer field specifies the client address lease time in seconds. This implementation of the DHCP server assigns 3 year leases by default, and the leases are expired only when explicitly surrendered by the client.

`pro_host_prefix`

This string field specifies the default text prefix for generating client hostnames. For example, the prefix `iris` directs `proclaim` to generate hostnames of the form `iris1`, `iris2`, `iris3`, and so on.

pro_choose_name

This boolean (true or false) flag specifies whether the client systems are allowed to choose their own hostname or whether they must be assigned the name given to them by the server. A value of 1 (true) in this field brings up a dialog box on the client system giving the user the option of either taking the name offered by the server or entering a hostname of the user's choice. If the user selects a name, the server allows this name if it passes basic tests for syntax and uniqueness, otherwise the server/client dialogue continues until a mutually acceptable name is submitted. A value of 0 (false) in this field indicates that the user on the client system must accept the name provided by the server. This option can be used with SGI DHCP clients and servers only. The client and server exchange DHCPREQUEST and DHCPOFFER packets until a mutually acceptable name is chosen. This option is being deprecated; use the `-H host-name` option with SGI clients to specify host names.

pro_ipaddress_range

This field takes an entry of integers using standard numeric range rules. The entry defines the range of host number addresses assignable by this server. For example, if the value is: 1-3, 5-7, 9. The server issues IP addresses with the base address specified in the configuration filename (such as *config.192.26.61.0*)

Each client is issued an IP address matching the name of the configuration file, suffixed with the numbers 1 through 3, and 5 through 7, and 9, but not 4 or 8. This option can be to restrict the IP addresses offered by a given server. This is very useful if the administrator wants to assign only certain block(s) of addresses using proclaim, or in the absence of a server to server protocol, wishes to have multiple servers serve clients on the same subnet.

pro_router_addr

This field of comma-separated IP addresses specifies a list of addresses for network routers on the client's subnet. Routers should be listed in the order of preference for their use.

pro_bootfile_size

Specifies length in 512-octet blocks of the default boot image for the client.

<code>pro_time_offset</code>	Specifies offset of the client's subnet from Coordinated Universal Time (UTC) in seconds.
<code>pro_meritdump_pathname</code>	Specifies a file for dumping a core image in case the client crashes.
<code>pro_root_pathnam</code>	Specifies the directory that contains the client's root disk partition
<code>pro_extensions_pathname</code>	Specifies a TFTP file containing vendor extensions for <i>bootp</i> . On SGI systems the this file has the same format as the <i>config.netnumber</i> file.
<code>pro_TFTPserver_name</code>	Specifies a TFTP server for the client.
<code>pro_bootfile_name</code>	Specifies the pathname of the default boot image for the client.
<code>pro_dns_domain</code>	This text field specifies the domain name that client should use when resolving hostnames using DNS.
<code>pro_dnsserver_addr</code>	This field of comma-separated IP addresses specifies a list of addresses for Domain Name System servers available to the client. Servers should be listed in the order of preference for their use.
<code>pro_propel_server</code>	This IP address specifies the address of the propel server that serves the clients on this subnet. Additional information about propel can be found in Chapter 2, "Software Distribution, User, and Host Management With propel," of the <i>IRIXpro Administrator's Guide</i> .
<code>pro_timeserver_add</code>	This field of comma-separated IP addresses specifies a list of addresses for time servers available to the client. Addresses should be listed in the order of preference for their use.
<code>pro_logserver_addr</code>	Specifies a list of UDP log servers available to the client.
<code>pro_cookieserver_addr</code>	Specifies a list of RFC 865 cookie servers available to the client.

- `pro_LPRserver_addr`
Specifies a list of RFC 1179 line printer servers available to the client.
- `pro_resourceserver_addr`
Specifies a list of RFC 887 resource location servers available to the client.
- `pro_swapservice_addr`
Specifies the IP address of the client's swap server.
- `pro_nis_domain`
This text field specifies the name of the client's NIS domain.
- `pro_nisserver_addr`
This field of comma-separated IP addresses specifies a list of addresses indicating NIS servers available to the client. Servers should be listed in the order of preference for their use.
- `pro_mtu`
This unsigned short integer field specifies the MTU (maximum transmission unit) to use on the network interface configured in this file. The minimum legal value for the MTU is 68.
- `pro_allnets_local`
This boolean (true/false) field specifies whether or not the client can assume that all other subnets of the IP network to which the client is connected use the same MTU as the subnet to which the client is directly connected. A value of 1 (true) indicates that all subnets share the same MTU. A value of 0 (false) means that the client should assume that some other subnets can have smaller MTUs.
- `pro_broadcast`
This IP address field specifies the broadcast address in use on the client's subnet.
- `pro_domask_disc`
This boolean (true/false) Perform Mask Discovery field specifies whether or not the client should perform subnet mask discovery using ICMP. A value of 1 (true) means that the client should perform mask discovery, while a value of 0 (false) indicates that the client should not perform mask discovery.

`pro_resp_mask_req`

This boolean (true/false) Mask Supplier field specifies whether or not the client should respond to subnet mask requests using ICMP. A value of 1 (true) means that the client should respond. A value of 0 (false) in this field means that the client should not respond.

`pro_static_routes`

This field takes a comma-separated list of routes in the following form: **dest_address - router_address, dest_address2 - router_address2**. The static route field specifies a list of static routes that the client should install in its routing cache. If multiple routes to the same destination are specified, they should be listed in descending order of priority. The routes consist of a list of IP address pairs. The first address is the destination address; its counterpart address, separated by a dash (-), is the address of the router to the destination. The default route (0.0.0.0) is an illegal destination for a static route.

`Serve_This_Network`

Toggle DHCP serving of this subnet.

Other Options

The proclaim server also handles other options listed in RFC 1533. To set values for these options use of the *ProclaimServerMgr*(1M) graphical user interface is recommended. You can view the standard config.Default file in the */var/dhcp/config* directory to see all the supported options.

The following are the active lines only from a working `config.netnumber` configuration file:

```
Serve_This_Network: 1
pro_address_counter: 1
pro_host_pfx_counter: 1
pro_netmask: 255.255.255.0
pro_lease: 604800
pro_host_prefix: n6634
pro_choose_name: 1
pro_ipaddress_range: 1-254
pro_dns_domain: engr.sgi.com
pro_nis_domain: engr.sgi.com
pro_mtu: 1500
pro_allnets_local: 1
pro_domask_disc: 0
pro_resp_mask_req: 0
```

For additional information about the `dhcp_bootp`, see the `dhcp_bootp(1M)` reference page.

Configuring the DHCP Relay Agent

The `dhcp_relay` program is a subnet agent that communicates with the main DHCP server to provide DHCP clients with host configuration information. If your site uses DNS to maintain the hosts map, you need some external mechanism (such as a script) to update the DNS map from the DHCP server. If your site uses NIS to maintain the hosts and ethers maps, your site's main DHCP server must be on the same machine as the NIS master. In any case, you should run `dhcp_relay` on all subnets besides the one with the main DHCP server.

The DHCP relay agent reads the configuration file `/var/dhcp/config/dhcp_relay.servers` to determine the location of the main DHCP server. This configuration file should contain the IP address and hostname of the DHCP server, on two separate lines.

Use the following command to invoke `ProclaimRelayMgr`.

```
/usr/sbin/ProclaimRelayMgr
```

The the `ProclaimRelayMgr` graphical interface (GUI) appears, as shown in Figure 3-5

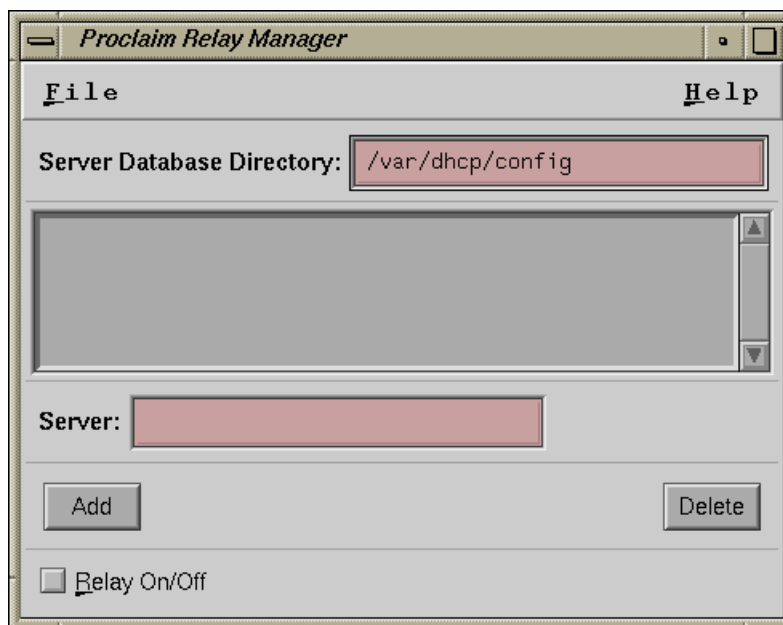


Figure 3-5 ProclaimRelayMgr Graphical Interface (GUI)

The main window allows you to add and delete addresses and names of proclaim servers. You can enable or disable the relay-agent by clicking the **Relay On/Off** button in the main window.

The *File* button opens a pull-down menu for saving the options and server addresses displayed in the main window to a file (using the *Save* button). To view messages logged by the relay-agent in the system log file (*/var/adm/SYSLOG*), launch *sysmon(1M)* from the pull-down menu by choosing the **System Log Viewer** item.

To control the execution of *dhcp_relay*, perform the following:

1. Specify relay-agent options in the `/etc/config/dhcp_relay.options` file, if any, as in this example: `-c proclaim_config_dir` This specifies an optional proclaim server configuration directory. The default directory is `/var/dhcp/config`.
2. The configuration directory contains a file `dhcp_relay.servers`, which has a list of proclaim servers, one per line. DHCP messages are forwarded to each server contained in this file. Proclaim servers may be specified using either an Internet host name or an Internet address in standard dot notation (see *inet(3N)*). Here is a sample `dhcp_relay.servers` configuration file:

```
192.26.61.9
dhcpserver.sgi.com
```
3. Enable the *proclaim* relay-agent by entering these commands:

```
chkconfig proclaim_relayagent on
chkconfig proclaim_server off
```
4. Replace the `bootp` entry in `/usr/etc/inetd.conf` with this line:

```
bootp dgram udp wait root /usr/etc/dhcp_relay dhcp_relay -o \
/etc/config/dhcp_relay.options
```
5. Kill any running `bootp` with this command:

```
/etc/killall bootp
```
6. Restart *inetd* with this command:

```
etc/killall -HUP inetd
```

For more information on how to configure the relay agent by running the *ProclaimRelayMgr* GUI, see the *ProclaimRelayMgr(1M)* reference page. If you prefer, you can configure relay agent options using your favorite text editor, following the steps documented on the *dhcp_relay(1M)* reference page.

Setting up the DHCP Proclaim Client

The *proclaim* client communicates with a DHCP server to obtain host configuration parameters, including an IP address and an IP address lease at minimum. You can use *proclaim* to set up and configure new systems automatically, and also to move systems from one network to another without administrative intervention. The *proclaim* client can also verify configurations at reboot time.

Only the superuser can employ *proclaim*. If the primary network interface is changed, *proclaim* updates the `netaddr` variable in the NVRAM.

In the absence of the `DHCPoptionsToGet` keyword, *proclaim* requests the following configuration parameters from an available DHCP server:

- an IP address
- the lease duration (3 year default)
- the subnet mask
- hostname (users may be allowed to choose their own)
- DNS domainname
- NIS domainname if the NIS option is installed.
- IP address of the DHCP server; see the `dhcp_bootp(1M)` man page.

If the server is configured in a manner where clients from specific subnets are allowed to choose their own hostname, then at the receipt of the `DHCPOFFER` message the client requests user input to either accept the offered hostname or specify a different one. This dialog between the server and client continues until a valid and unique mutually acceptable name is selected. Alternately, the `-w` option can be used to request a hostname and/or IP address non-interactively. The client will print an error and exit if the requested hostname or address is invalid.

Proclaim sleeps after obtaining a new lease or after verifying an existing lease, until it is time to renew the lease. The client wakes up at times specified by the DHCP protocol and attempts to renew the lease until it succeeds or the current lease expires. If the client is unable to renew the lease then it can shutoff network access depending on the use of the `-E` option.

To set up a the DHCP Proclaim Client, follow these steps:

1. Edit the `/etc/config/proclaim.options` file if necessary.
2. Edit the `/usr/etc/dhccopt.local` if you want to add to the processing of options
3. Execute the following: `chkconfig autoconfig_ipaddress on`
4. Optionally, backup files which may be changed by the `dhccopt` script as follows:

```
cp /etc/resolv.conf /etc/resolv.conf.orig
cp /etc/config/static-route.options
/etc/config/static-route.options.orig
```

5. Determine whether you need to change the `/etc/hosts` or `/etc/sys_id` files as follows:
 - a. IRIS is the host name by default, if this is acceptable you do **not** have to change the files
 - b. If you have configured your system to have the DHCP server supply the host name, no change is necessary.
 - c. If neither a nor b applies, change the `/etc/sys_id` and `/etc/hosts` files to reflect the desired host name.
6. Reboot the system

The command is as follows:

```
/usr/etc/proclaim [ -B ] [ -b lease_time ] [ -d -s server_addr ] [ -E ] [ -e 2nd_interface ] [ -I ] [ -i ] [ -l lease_time ] [ -p ] [ -r lease_time -s server_addr ] [ -t repeat ] [ -x max_timeout ] [ -H hostname ] [ -w -h host_name -a ip_addr ]
```

The `proclaim(1M)` command accepts the following options:

- B The **-B** option broadcasts a BOOTP packet to a BOOTP 1533 server. This is useful when you just want to obtain the other configuration data, and the IP address is not otherwise known to the BOOTP or DHCP server. The data is placed in the file `/var/adm/proclaim.data`. `proclaim` does not change any networking state with this option.
- b The **-b** *lease_time* option is used to manually rebind the lease; a *lease_time* in seconds must be specified.
- d The **-d** option allows the client to surrender its address lease and shutdown all networking. This option can only be used in conjunction with the **-s** *server_address* option to specify the Proclaim server address. This usually is the server from where the original address lease was obtained. If there is a client running on the host use the **`/etc/init.d/run-proclaim surrender`** command to surrender the lease.
- E The **-E** option is used to specify that network access should be shut off if the client is unable to renew the lease and the lease expires. The default behavior does not shut off network access if the lease expires.
- e The **-e** *2nd_interface* option is used to specify the device for a second network interface.

- `-I` The `-I` option is used to send a DHCPINFORM packet to a DHCP server. This is useful when an IP address was assigned to the host using some other means and it is required to obtain only other configuration parameter.
- `-i` The `-i` option is used where the client is seeking to verify a previously allocated, cached configuration. If the client receives a DHCPNAK response to this message, it means that either the 'requested IP address' parameter is incorrect or the client is on the wrong network. The client automatically requests a new address and lease and proceeds to install the new configuration. If there is no response, it means the server(s) has no record of this client. These messages are logged using the system logging daemon `syslogd(1M)`. If a client is running on the host use the `/etc/init.d/run-proclaim` verify command to verify the lease.
- `-l lease_time` The `-l lease_time` option is used to request a different address lease duration than the default of 3 years. The lease time needs to be specified in number of seconds.
- `-p` The `-p` option is used to print the status of the current lease and other configuration parameters. This status is obtained from the `/var/adm/proclaim.lease_info` file created at the time the lease was obtained.
- `-r lease_time` See the `-s server_addr` option.
- `-s server_addr` The `-r lease_time -s server_addr` option is used to manually renew the lease; a `lease_time` in seconds must be specified. The `server_addr` specifies the server with which to renew the lease. The server address may be obtained by using the `-p` option of `proclaim`.
- `-t invocations` The `-t invocations` option is used to specify the number of invocations of the client using the `run-proclaim(1M)` script, after which the `autoconfig_ipaddress` flag is set to off (see `chkconfig(1M)`). When this option is not specified the client may be run up to a default 2 (or as specified in the `/etc/config/proclaim.options` file) number of times using the script after which the flag is set to off. Use of the script to run the client requires the flag to be set to on.
- `-x max_timeout` The `-x max_timeout` option is used to specify a maximum timeout (time to wait for a server to reply) in seconds. The current default is set to 6 seconds.

- H *hostname*** The **-H *hostname*** option is used to specify a hostname using the DHCP Host Name Option. The `run-proclaim(1M)` script uses this option to request the current hostname of the client from the server.
- w -h *host_name*** The **-w -h *host_name*** option **-a *ip_addr*** option is used to choose a hostname and IP address without the user being prompted by the client. If no *host_name* or *ip_addr* is specified, the the hostname and IP address offered by the server will be used. If only the *host_name* is specified, the server will choose an appropriate IP address. This option was added to make it possible to run proclaim from a CGI script

Options File

Options may also be specified in the `/etc/config/proclaim.options` file. Options specified on the command line supersede those specified in this file. Lines beginning with a `#` are treated as comments. The option keywords in this file must be followed immediately by a colon, then by any number of tabs or spaces, and finally by the value of the option

The following keywords are supported:

- `Invocations` - equivalent to **-t** option
- `MaxTimeout` - equivalent to **-x** option
- `ServerAddress` - equivalent to **-s** option
- `ShutdownOnExpiry` - equivalent to **-E** option
- `Lease` - equivalent to **-l** option
- `DHCPoptionsToGet` - additional options to request from DHCP server
- `DHCPoptionsRenew` - options to request for lease renewal/rebind

The values for `DHCPoptionsToGet` and `DHCPoptionsRenew` are specified by listing the option numbers separated only by a comma. For example,

```
DHCPoptionsToGet:      3,6,26,27,41
```

`DHCPoptionsToGet` lists options to get in addition to the standard (default) configuration. This affects the initial lease request.

For a complete list of options, see RFC 2132.

DHCPoptionsRenew specifies what options should be included in the renew/rebind request. DHCPoptionsRenew is expected to be a subset of the options requested when the lease is acquired (the standard configuration parameters plus the additional ones listed in DHCPoptionsToGet).

Options returned from a renew/rebind request will be checked. If basic values have changed (IP address, host name, subnet mask, DNS domain, NIS domain) the host will be reconfigured and the /usr/etc/dhccopt script will be run. If only non-basic values have changed, only the script will be run. The dhccopt script file contains processing of the gateway, DNS domain and NIS domain. A script file /usr/etc/dhccopt.local may be created to add processing of options. If this file exists, dhccopt calls it for additional actions for the options returned.

Additional options which are not keywords specify whether the run-proclaim script should try to obtain configuration parameters for all the network interfaces on the system. If none of these options are present then the proclaim client is invoked for the primary interface only.

The options are:

```
Primary on|off
Interface
```

The presence of an on invokes the proclaim client and off does not invoke it for the specified interface when the run-proclaim script is executed with a start argument. The value of interface is the name, for example, ecl, xpi1, etc.

For more information about setting up *proclaim* on client systems, see the *proclaim(1M)* reference page.

Limitations of DHCP

The following restrictions and limitations are present in this release:

- The DHCP server must be the NIS master if your site uses NIS for hostname, IP address, or network hardware address validation and mapping. Note that NIS is an optional software product, and not all systems and networks use it.
- If your site uses DNS to maintain the hosts map, you need an external mechanism (manual or automated) to update the DNS map from DHCP. This limitation should be lifted in the next release.

- The server cannot respond to both DHCP and *bootp* client requests contained within a single packet, although it can act as the *bootp* server for normal *bootp* requests.

Creating a Local Network Script

To start and terminate locally developed network daemons, or to publish ARP entries and set routes, create the separate shell script `/etc/init.d/network.local`. Make symbolic links in `/etc/rc0.d` and `/etc/rc2.d` to this file to have it called during station startup and shutdown:

```
ln -s /etc/init.d/network.local /etc/rc0.d/K39network
ln -s /etc/init.d/network.local /etc/rc2.d/S31network
```

See `/etc/init.d/network` for the basic format of the script. Also refer to `network(1M)`, `rc2(1M)`, and `rc0(1M)` for more information.

Turning On Remote Access Logging

Several network daemons have an option that lets you log remote accesses to the station log file `/var/adm/SYSLOG` by using `syslogd`. Sites connected to the Internet should use this feature. To enable logging for `ftpd`, `tftpd`, and `rshd`, edit `/etc/inetd.conf` and add `-l` after the right-most instance of `ftpd` and `tftpd`. Add `-L` after the right-most instance of `rshd`. For additional `ftp` logging, add `-ll` to the `ftpd` entry. Signal `inetd` to reread its file after you have added your changes:

```
/etc/killall -HUP inetd
```

Remote logins by means of `rlogin`, `telnet`, and the 4DDN `sethost` programs can be logged by `login`. Edit `/etc/default/login` and add the keywords `SYSLOG=ALL` or `SYSLOG=FAIL` to it. For example, this command in the `login` file logs successful and failed local and remote login attempts to `syslogd`:

```
syslog=all
```

See the `login(1)` reference page for details.

Setting Up Network-Wide Services

In addition to the required software that you must set up on each system on the network, you may want to designate certain systems to provide network-wide services. For example, you can set up a system as an InfoSearch server, so that users can access a complete set of InfoSearch books without installing them all on their local disks.

The following sections discuss setting up an anonymous *ftp* server, and setting up an InfoSearch server.

- “Setting Up an Anonymous FTP Account” on page 92
- “Setting Up a Password-Protected FTP Account” on page 97
- “Setting Up a Conventional InfoSearch Server/Client System” on page 99
- “Setting Up a CD-ROM InfoSearch Server/Client System” on page 101
- “Starting Up the Remote InfoSearch Viewer” on page 102

Setting Up an Anonymous FTP Account

An anonymous FTP account is a way for you to make information on your system available to anybody, while still restricting access to your system. An anonymous FTP account lets anybody log in to your system as the “anonymous” or “ftp” user. The FTP daemon does a *chroot* to the FTP home directory (*~ftp*) for anonymous FTP users, which effectively prevents them from accessing parts of the system that are not subdirectories of *~ftp*; see *chroot(1M)*. If you prefer to limit access to users with passwords, after setting up an anonymous account continue with “Setting Up a Password-Protected FTP Account” on page 97.

The following procedure is designed to help you set up a network-accessible anonymous FTP account. As usual, understanding the various steps and continually monitoring how the account is used are necessary to protect your system security.

Note: If you are using IRIX 6.3, 6.4, or 6.5, you must copy certain files into the */lib32* and */usr* directories. See “Updating FTP for IRIX 6.3, 6.4, and 6.5” on page 98 for details. Otherwise follow the steps below.

1. Create the anonymous FTP user entry in `/etc/passwd`. The username should be `ftp`. Put an asterisk (*) in the password field, and assign user and group IDs, a home directory, and a login shell. The following is an example of a typical entry in `/etc/passwd` for the anonymous FTP account:

```
ftp:*:997:995:Anonymous FTP Account:/disk2/ftp:/dev/null
```

The login shell `/dev/null` is recommended but not required, and the home directory can be anywhere, with reservations as explained in the next step.

Tip: For public servers, SGI recommends that you create a shadow password file. Run the `pwconv` command from the `/etc` directory:

```
# pwconv
```

This command updates the contents of `/etc/passwd` and moves encrypted passwords to `/etc/shadow`, which is then inaccessible to non-privileged users. If you run NIS, see the `shadow(4)` reference page for interoperability information.

2. Create an FTP home directory. This may be wherever you like but, especially if you are going to allow writes to it, it should probably be on a separate partition from `/` or `/usr`. That way, if the partition fills up, it will not disable basic system operations. In this example procedure, `/disk2/ftp` is the name of the anonymous FTP directory. First, make the directory:

```
# mkdir /disk2/ftp
```

Then, if it is a separate disk or disk partition, you can mount the device on it; see `mount(1M)`. The anonymous FTP home directory you make must be the same one you specify in the `/etc/passwd` file.

3. Change directory to the `ftp` home directory and create the subdirectories used for FTP access:

```
# cd /disk2/ftp
# mkdir bin dev etc lib pub incoming
```

In addition to the standard `bin`, `dev`, `etc`, `lib`, and `pub` directories, you may wish to create an `incoming` directory for incoming files.

4. Copy the `ls` command from `/sbin` to `~ftp/bin`:

```
# cp /sbin/ls bin
```

Note: In order for the *ls* command to work with an IRIX 6.5.x system, perform the following steps:

```
# mkdir /dir2/ftp/lib32
# cp /lib32/libc.so.1 /dir2/ftp/lib32
# cp /lib32/rld /dir2/ftp/lib32
# chmod -R 555 /dir2/ftp/lib32
```

5. Copy */etc/passwd* and */etc/group* to *~ftp/etc* and edit them to an acceptable minimum:

```
# cp /etc/passwd /etc/group etc
```

A good choice for the contents of *~ftp/etc/passwd* might be

```
root:*:0:0:Super-User:/:/dev/null
bin:*:2:2:System Tools Owner:/bin:/dev/null
sys:*:4:0:System Activity Owner:/var/adm:/dev/null
ftp:*:997:999:Anonymous FTP Account:/disk2/ftp:/dev/null
```

A good choice for the contents of *~ftp/etc/group* might be

```
sys:*:0:
other:*:995:
guest:*:998:
```

6. Add appropriate device and library files for anonymous FTP as follows:

```
# /sbin/mknod dev/zero c 37 0
# cp /lib/libc.so.1 /lib/rld lib
```

The *dev/zero* file helps zero out sensitive data; *~ftp/bin/ls* requires the library files.

7. Set restrictive permissions on *~ftp/etc/passwd*, *~ftp/etc/group* and *~ftp/dev/zero*:

```
# chmod 444 etc/* dev/*
```

8. Make sure the *bin*, *dev*, *etc*, *lib* and *~ftp* directories are owned by root, with group being *sys*, having restricted write permission:

```
# chown root.sys bin dev etc lib lib32 .
# chmod 555 bin dev etc lib lib32 .
```

In the *chown* command, the dot separates owner and group.

9. For the *pub* directory, set the owner to root, the group to *sys*, and global read and access permission:

```
# chown root.sys pub
# chmod 755 pub
```

10. If you created an *incoming* directory, set the permissions to allow anybody to write there but not to read its contents:

```
# chown ftp.other incoming
# chmod 333 incoming
```

Any FTP user can now get or put files in the *incoming* directory, but they must know the name of the file beforehand, because they cannot list directory contents.

Caution: By allowing write permission, you make it possible for anonymous FTP users to fill the entire disk partition. For instructions on setting up a more restricted alternative, see “Setting Up a Password-Protected FTP Account” on page 97.

11. As a security precaution, add the following entry to the */etc/aliases* file to cause mail sent to the user *ftp* to be intercepted by the postmaster:

```
ftp: postmaster
```

Run the command *newaliases* to make this take effect. (This assumes you have an alias of postmaster in */etc/aliases*. See *aliases(4)* and *newaliases(1M)*.)

12. Enable FTP logging as described in “Limiting inetd Services” in Chapter 5 in *IRIX Admin: Backup, Security, and Accounting*. Note that if you use one *-l* argument with *ftpd*, you record only successful and failed FTP login attempts. If you use two *ls*, you also record the retrieve (*get*), store (*put*), append, delete, make directory, remove directory, and rename operations (and their filename arguments) performed during ftp login sessions. If you use three *ls*, the report includes the number of bytes transferred in *get* and *put* operations.

For example, the following entry in */etc/inetd.conf* means FTP sessions and get or put operations (excluding byte count) are logged in */var/adm/SYSLOG*:

```
ftp      stream  tcp      nowait  root    /usr/etc/ftpd  ftpd -ll
```

13. Once you have edited */etc/inetd.conf*, restart *inetd* with the following command:

```
# /etc/killall -HUP inetd
```

Note: Although FTP logging records in */var/adm/SYSLOG* show any passwords entered by users logging in, no password checking is done for anonymous FTP. The convention is for anonymous users to enter their e-mail addresses for passwords, but they could just as easily enter another user's address or anything at all.

If you place text in the */etc/issue* file, it displays when a user connects to your system, before the FTP login prompt. You might want to include information here about the kind of services your FTP site offers, and whom to contact in case of problems. In addition, text in the *~/ftp/README* file displays after an FTP user logs in.

Refer to *crontab(1)*, *syslogd(1M)*, and the file */var/spool/cron/crontabs/root* for information on changing the frequency or nature of system log file maintenance—you may, for example, want to increase the length of time you keep log files. To help you keep track of the demands made on your public FTP server, see the IRIX Admin: Backup, Security, and Accounting book. Chapter 6 provides information on auditing system resource usage, and Chapter 7 provides general system accounting information.

Setting Up a Password-Protected FTP Account

Password-protected FTP accounts allow only valid users to access your system. This is a way to make file transfer available to selected people, while still restricting access to the rest of your system. Only users with a password in */etc/passwd* may connect to these FTP services. Additionally, restricted FTP users cannot access any part of a system that is not part of their home directory; the system does this with *chroot*.

The following procedure is designed to help you set up a network-accessible secure FTP account. As usual, understanding the various steps and continually monitoring how the account is used are necessary to protect system security. Before beginning this procedure, you must have completed all steps in “Setting Up an Anonymous FTP Account” on page 92.

1. Modify the `ftp` user entry in */etc/passwd*. The username should be “ftpX” where X is a variable. Leave the password field empty for now, assign a user ID and group ID, and home directory. The following is an example of a typical entry in */etc/passwd* for a password-protected FTP account:

```
ftpl: :197:995:FTP user:/disk2/ftp:/dev/null
```

The login shell */dev/null* prevents FTP users from gaining UNIX shell access. The second field will contain the user’s encrypted password.

Tip: The FTP daemon grants anonymous FTP privileges to the user named `ftp`, so if there is no `ftp` line in */etc/passwd*, the system does not permit anonymous FTP.

2. Assign the FTP user a reasonable password, maybe asking what the person prefers:

```
# passwd ftpl
New password:
Re-enter new password:
```

3. Copy the line for the FTP user from */etc/passwd* and paste it into *~ftp/etc/passwd*, changing the encrypted password to an asterisk:

```
ftpl:*:197:995:FTP user:/disk2/ftp:/dev/null
```

4. Create or edit the */etc/ftpusers* file so it contains this line:

```
ftpl restrict
```

This causes the FTP daemon to *chroot* this user to its home directory; see `ftpd(1M)`.

5. Create the FTP user's home directory inside *ftp/pub*:

```
# mkdir /disk2/ftp/pub/ftpl
```

6. For the FTP user's home directory, set ownership and group *other*, and give the user read, write, and access permissions:

```
# chown ftpl.other /disk2/ftp/pub/ftpl
# chmod 700 /disk2/ftp/pub/ftpl
```

Caution: With write permission, FTP users can fill up the containing disk partition, just as was true with the *~ftp/incoming* directory.

7. Create the file *~ftp/README* telling FTP users where to find their home directory:

```
Your home directory is under pub/.
```

8. As a security precaution, add the following entry to the file */etc/aliases* to cause mail sent to the FTP user to go to the postmaster instead:

```
ftpl: postmaster
```

Run the command *newaliases* to make this take effect; see *newaliases(1M)*.

Updating FTP for IRIX 6.3, 6.4, and 6.5

If you are running IRIX 6.3, 6.4, or 6.5, you will need to copy the runtime loader into the */lib32* directory, as well as */usr/lib/iconv/iconvtab* into the */usr/iconv* directory.

1. Create a directory for *~ftp/lib32*:

```
# mkdir ~ftp/lib32
```

2. Copy the runtime loader into the directory:

```
# cp /lib32/rld ~ftp/lib32
```

3. Create a directory for *~ftp/usr/iconv*:

```
# mkdir ~ftp/usr/iconv
```

4. Copy the font files into the directory:

```
# cp /usr/lib/iconv/iconvtab ~ftp/usr/lib/iconv
```

About InfoSearch File Servers

The files and directories that make up the InfoSearch system of online documentation take up a great deal of space. In your network, there is no reason for each system to maintain a separate copy of the InfoSearch documents as long as all systems are at substantially the same software revision level. If the InfoSearch software revision level is the same across several systems, you can designate one system to serve the others with the InfoSearch files, thus freeing up disk space on the client systems.

Be sure to choose a server system that is not called upon to carry a heavy workload, and take your network performance into account before you begin. If your users use InfoSearch a great deal and your network is already burdened, you may find that your users will not appreciate the decreased response time from both InfoSearch and your network in general. Also, if a person is to use the InfoSearch server as his or her workstation, that person must be prepared to accept a certain (possibly substantial) amount of disk, CPU, and network overhead as a result.

There are two convenient ways to set up an InfoSearch server. These will be detailed in the next two subsections. Both methods require that you have the NFS software option installed. If you do not understand the terms and concepts behind NFS, you should read the *ONC3/NFS Administrator's Guide* and the *NIS Administration Guide* before undertaking these projects. The second method described here also requires a dedicated CD-ROM drive to hold the InfoSearch distribution media.

Setting Up a Conventional InfoSearch Server/Client System

To install the InfoSearch Viewer and Document Library on a remote server and retrieve the information from a local client system, follow the steps below.

On the server system:

1. Log in as *root* (superuser).
2. Bring up *inst* and install the complete InfoSearch product image (from your CD-ROM drive or distribution directory) with these commands:

```
Inst> install InfoSearch InfoSearch_gloss *.books.*  
Inst> go
```

The total size of the viewer and document library is a little less than 23 megabytes.

3. Export the `/usr/share/InfoSearch/library/SGI_bookshelves` directory using the System Manager or the `exportfs` command:

```
exportfs -i /usr/share/InfoSearch/library/SGI_bookshelves
```

If the server does not have a graphical display, a warning is generated during the exit. This warning relates to updating the X server's font directory and can be ignored.

On the client system:

1. Log in as `root` (superuser).

2. Enter the command

```
versions remove *.books.*
```

to remove the books on your bookshelves.

3. Run the `inst` command on your client system and enter the following commands to install the `InfoSearch.sw.client` subsystem (you may want to install the `InfoSearch.man.man` and the `InfoSearch.man.relnotes` subsystems as well):

```
Inst> keep InfoSearch InfoSearch_gloss
```

```
Inst> install InfoSearch.sw.client
```

```
Inst> go
```

4. Mount the `SGI_bookshelves` directory from your server on your client system. In the example below, the name of the server machine is `capra`.

```
mkdir /usr/share/InfoSearch/library/SGI_bookshelves
```

```
mount capra:/usr/share/InfoSearch/library/SGI_bookshelves  
/usr/share/InfoSearch/library/SGI_bookshelves
```

Note that when you enter the `mount` command on your client system, the entire command line goes on a single line. The command is broken across two lines in this example due to formatting limitations.

Note: If you have remote-mounted your InfoSearch books, the SGI `desktophelp` system will not operate correctly while the books are mounted. To view the `desktophelp`, unmount the books temporarily.

Setting Up a CD-ROM InfoSearch Server/Client System

With this method, you need not use up your disk space for the InfoSearch files if you have a CD-ROM drive you can dedicate to InfoSearch. You simply leave the CD with the InfoSearch distribution in the drive and link the mounted distribution to the directory where InfoSearch would have installed the files. The drawback of this system is that you must dedicate a CD-ROM drive to the purpose, but this method can be used with NFS to provide server access to your entire network.

Note: If you have a version of the InfoSearch Document Library installed on your client disk but you want to mount the books from the CD-ROM, you must remove all the files in `/usr/share/InfoSearch/library` before creating the symbolic link described in step 2. Use the `versions remove` command to cleanly remove the books, then check the directory to be sure all files have been removed.

If you wish to use the InfoSearch Viewer and Document Library from the CD-ROM and access the information from a local or remote system, follow the steps below.

On the server system with the CD-ROM drive:

1. Log in as *root* (superuser).
2. Insert the InfoSearch CD into the CD-ROM drive. If the drive is not mounted, use the System Manager or the `mount` command to mount the drive. The most common mount point is `/CDROM`. If the drive is mounted correctly, you should be able to change directories to `/CDROM` and see all the files in the InfoSearch CD. To see the files, use the command
3. As *root*, create a symbolic link from

```
cd /CDROM/InfoSearch
```

`/usr/InfoSearch/library/SGI_bookshelves` to the CD `InfoSearch` directory. You may need to create the directory `/usr/InfoSearch/library` if it doesn't exist. Use these commands:

```
mkdir -p /usr/share/InfoSearch/library  
ln -s /CDROM/InfoSearch/SGI_bookshelves  
/usr/share/InfoSearch/library
```

Note that when you enter the link command on your client system, the entire command line goes on a single line. The command is broken across two lines in this example due to formatting limitations.

At this point, the InfoSearch bookshelves are mounted on your server and available for use on that system. To allow users on other systems on your network to use the bookshelves, proceed to step 4.

4. If you want to share the bookshelves with other users on your network, export the `/CDROM` directory using the System Manager or the `exportfs` command:

```
exportfs -i /CDROM
```

On each client system, perform the following steps:

1. Log in as `root` (superuser).
2. Run the `inst` command and give the following commands to install the `InfoSearch.sw.client` subsystem (you may want to install the `InfoSearch.man.man` and the `InfoSearch.man.relnotes` subsystems as well). You can get the installation software from the CD in the `/CDROM/dist` directory on the server.

```
Inst> keep *  
Inst> install InfoSearch.sw.client  
Inst> go
```

3. Mount the bookshelves from the server. In the example below, the name of the server machine is `capra`. Use the following command:

```
mount capra:/CDROM/InfoSearch/SGI_bookshelves  
/usr/share/InfoSearch/library/SGI_bookshelves
```

Note: When you enter the `mount` command on your client system, type the entire command on a single line. The command in the previous example is broken across two lines due to formatting limitations.

Starting Up the Remote InfoSearch Viewer

After the software has been installed, you should force your shell to remake its list of available programs and commands with the command `rehash`.

You can now invoke the InfoSearch Viewer. The command to invoke the viewer is `iiv` (an acronym for InfoSearch Viewer) `iiv`.

Setting up Time Synchronization Service

Table 3-2 shows the three general-purpose IRIX time synchronization services.

Table 3-2 IRIX Time Synchronization Services

	Subsystem	Installed by Default	On by Default
ntp(1M)	ntp	no	no
timed	eeo.sw.base	yes	yes
timeslave	eeo.sw.base	yes	no

Note that the network time protocol (NTP) is a newly supported component as of the IRIX 6.5.22 release.

Table 3-2 shows that the `timed` service is **on by default** and is configured to synchronize the clock. In general, only **one** service should be configured on in order to prevent any dispute and confusion among the time service daemons.

Table 3-3 shows the default configuration parameters used for each time service.

Table 3-3 IRIX Time Synchronization Services Default Configuration Parameters

	Parameter file in /etc/config	Default Parameters
ntp(1)	ntp.options	<code>-b -f /etc/ntp.drift -c /etc/ntp.conf</code>
timed(1M)	timed.options	<code>-G timelords -P /var/adm/timetrin</code>
timeslave(1M)	timeslave.options	No file is supplied.

Several time synchronization configuration scenarios are possible, in particular, using `timeslave` as a slave to a NTP server and using `timed` as a time server for a local subnet. Such configuration should be done with care. For more information, see the `ntp(1M)`, `timed(1M)`, and `timeslave(1M)` man pages.

Note: The `ntp` man pages are HTML documents in `/usr/share/doc/ntp`.

For more information on NTP, see the following man pages and documentation:

- General NTP Overview, configuration
file:/usr/share/doc/ntp/index.htm
- ntp.conf - ntp configuration file
file:/usr/share/doc/ntp/confopt.htm
- ntpd - Network Time Protocol (NTP) daemon
file:/usr/share/doc/ntp/ntpd.htm
- ntpq - standard NTP query program
file:/usr/share/doc/ntp/ntpq.htm
- ntpdc - special NTP query program
file:/usr/share/doc/ntp/ntpdc.htm
- ntpdate - set the date and time via NTP
file:/usr/share/doc/ntp/ntpdate.htm
- ntptrace - trace a chain of NTP servers back to the primary source
file:/usr/share/doc/ntp/ntptrace.htm

Accessing Network Services through Internet Gateway

Beginning with IRIX 6.5, you can configure access to the Internet using the Internet Gateway software, which allows access to network services using a graphical user interface. Online help provides guidance; this section outlines the processes.

The interface allows you to administer FTP, DHCP, various naming services, and routing from the Network Services entry screen, as well as connecting to the Internet through an Internet Service Provider.

To access the Internet Gateway server, bring up a Web browser, and fill out a series of forms that complete the configuration process.

1. Ensure the server is running, then use the *chkconfig* command to turn the Internet Gateway on:

```
# chkconfig webface on
```

2. Start the Internet Gateway web server:

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

The Web browser can be run from any computer or workstation that is on the same LAN as the Internet Gateway server.

3. Open the following uniform resource locator (URL):

```
http://server_ip_address:2077/
```

In this URL, *server_ip_address* is the IP address you assigned to the Internet Gateway server. For example:

```
http://151.166.96.36:2077/
```

You will be prompted to enter an administrative account user ID and create a password for the administrative account. You will also be required to enter the root password.

You can now follow the online instructions to continue setting up your server to act as an Internet gateway.

Reserving Resources With RSVP

Beginning with IRIX 6.5, the Resource Reservation Protocol (RSVP) allows applications to set up pathways through computer networks by reserving bandwidth. This is particularly useful for applications such as video conferencing, video broadcasting, and Internet telephony, that may wish to guarantee a specific bandwidth, regardless of the current network load.

Of the two types of integrated services defined for RSVP, controlled load and guaranteed, only the controlled load service is supported. This service ensures that packets are delivered with minimal loss and delay. Each host and router along the pathway has the option of committing to provide the service or not, depending on current commitments and other factors.

An application communicates with the IRIX system daemon *rsvpd*, using the API defined by X/Open. The daemon *rsvpd* then sends and receives RSVP control messages on behalf of the application. *rsvpd* sends PATH messages advertising the availability of one or more of the integrated services, and those systems receiving a PATH message may send back a RESV message to reserve bandwidth. A packet handler in the IRIX kernel identifies packets belonging to an RSVP session, and gives them the quality of service

defined by the controlled load specification. The packet handler does this while maintaining high line utilization and low CPU overhead.

Installing RSVP

Before attempting to use RSVP, make sure a supported network interface together with the appropriate driver are installed on the system. To determine which network interfaces are present on the system, use the `hinv` command:

```
% hinv
```

The output should reveal the output as listed in Table 3-4.

Table 3-4 Supported Cards for RSVP

Type of Network Interface	Designation	Output from hinv (where <i>n</i> is an integer)
FDDI	xpi	xpin
FDDI	ipg	ipgn
FDDI	rns	rnsn
Ethernet	me	ecn
Ethernet	ec2	ecn
Ethernet	ef	efn
Ethernet	ecf	ecfn
Ethernet	ee	etn
Ethernet	ep	epn

With the supported hardware installed, verify that the appropriate level of the IRIX operating system is installed. You can check currently installed software using the desktop Software Manager, or from the command line, using `showprods`. If you need to install the software, refer to *IRIX Admin: Software Installation and Licensing* for instructions.

When your system is configured to handle RSVP, the next step is to start the RSVP daemon, */usr/etc/rsvpd*. To have *rsvpd* start automatically during system boot, use the *chkconfig* command:

```
# chkconfig rsvpd on
```

This should then activate your system to use the RSVP facility next time you boot. The *rsvpd(1M)* reference page provides details of available command line options.

The state of network bandwidth reservation may be monitored using the command *rstat*. The output of the command identifies each interface, line by line, the session address and port, and the next hop address. If no reservations exist for the session, *no resv* is displayed. See the *rstat(1M)* reference page for details and options.

RSVP can be configured on a particular interface using the command *psifconfig*. You can set the reservable bandwidth on the interface, see how much bandwidth is currently reserved, or completely disable or restore RSVP on that interface. See the *psifconfig(1M)* reference page for details of options.

If you prefer, you can configure *rsvpd* and monitor current RSVP status in graphical form using the Internet Gateway, see “Accessing Network Services through Internet Gateway” on page 104. The following options may be handled in this way:

- RSVP configured on or off with *chkconfig*.
- Determine which sessions are being advertised and reserved with *rstat*.
- Configure RSVP on a particular interface with *psifconfig*.
- Logging level may be set to control the logging of data for debugging.

You can, however, use RSVP from the command line, without the Internet Gateway.

Troubleshooting RSVP

Because RSVP is an extremely new technology, some troubleshooting should be expected. The following represent the first effort at determining what may occur.

Slow File Transfer

If you experience a marked slowdown in network file transfer, RSVP may be interfering with normal transmission. Your first action should be to turn off RSVP, using the kill command:

```
# killall -TERM rsvpd
```

Once you have benchmarked file transfer, restart RSVP.

Video Interruption

Smooth, seamless delivery of video picture should be the norm using RSVP, and a RSVP-capable video application. If picture delivery becomes jerky or interrupted, run *rstat* to see what reservations are in place on that machine. See the *rstat(1M)* reference page for explanation of output.

Troubleshooting Your Ethernet Connection

This section addresses some of the common Ethernet-related problems:

- “Troubleshooting Cable Problems” on page 108
- “Troubleshooting Late-Collision Problems” on page 109
- “Troubleshooting Packet Size Problems” on page 110
- “Troubleshooting Server Contact Problems” on page 110

See also “Troubleshooting System Configuration Using System Error Messages” in *IRIX Admin: System Configuration and Operation*.

Troubleshooting Cable Problems

Unplugged or loose cables are the most common problem that causes Ethernet-related error messages. For example:

```
unix: <ethernet_device>: no carrier: check Ethernet cable
```

This message means that the carrier was not detected when attempting to transmit. One recommendation is to check to see if the Ethernet cable is unplugged from the back of the

machine. For detailed instructions on connecting the cable, see your owner's guide. You do not need to shut down the system to connect or disconnect the cable. If you re-connected the cable, test the network connection.

After determining that the transceiver cable is plugged into the back of the machine and also into the transceiver box, look for other possible reasons for failure.:

- Check the machine's transceiver.
- Check the transceiver cable, and try swapping it out with one that is working on another machine.
- Check for a problem with a 10baseT hub.
 - Try a different port on the hub.
 - Try a different cable.
 - Try a different hub.
- If you try all the troubleshooting techniques and you still cannot access the network, contact your service provider; the network itself may be temporarily out of order.

For more information on this error, see the reference page for ethernet(7).

Troubleshooting Late-Collision Problems

Another problem related to cabling is that of late collisions. These result in errors shown in syslog detailing an error about late collisions.

Often the controller tried to transmit a packet but received a late collision signal from another machine. This usually indicates a problem in the Ethernet cable layout. The most frequent causes of this problem are excessively long cables and loose cable connections.

To isolate the problem, first look for violations of the Ethernet cable limits. For 10BASE5 cable (thicknet) the maximum segment length is 500 m; for 10BASE2 cable (thinnet) the maximum segment length is 200 m. The length of a transceiver cable (the cable that goes between a machine's Ethernet port and its transceiver) should not exceed 50 meters. If all cables are within the specified limits and all connectors are firmly seated, this error could indicate a hardware problem in an Ethernet controller or transceiver. Some recommendations are:

- Check other machines on the network for bad Ethernet controllers.

- Check other machines on the network for bad transceivers. If you suspect a bad transceiver, you can try swapping in a different one that works on another system.
- Check to see if removing a certain machine from the network makes the problem change or disappear.

Troubleshooting Packet Size Problems

A bad Ethernet controller somewhere on the network can cause problems by sending packets that are too large or too small, causing errors like this:

```
unix: <ethernet_device>: packet too small (length = <packet size>)  
unix: <ethernet_device>: packet too large (length = <packet size>)
```

The Ethernet controller received a packet that was smaller than the minimum packet size or larger than the maximum packet size allowed by Ethernet. This problem is caused by another machine with a bad Ethernet controller or transceiver. Some recommendations are

- Check other machines on the network for bad Ethernet controllers.
- Check other machines on the network for bad transceivers. If you suspect a bad transceiver, you can try swapping in a different one that works on another system.
- Check to see if removing a certain machine from the network makes the problem change or disappear.

For more information on this error, see the reference page for ethernet(7).

Troubleshooting Server Contact Problems

When your system cannot contact a network system, an error message similar to this is displayed:

```
portmapper not responding: giving up
```

This problem occurs in one of these situations:

- The system is not running.

Physically go to the system or call the system's Primary User or Administrator to check to see if the system is powered on and running.

- The network is not running.
To check, try to access another system on the network.
- The network administrator changed some information about the system or about the system's logical location on the network.
Check whether this is the case, and get the appropriate information to fix the problem.
- The system has too many users or systems using its resources. It cannot provide services to you at this time.
Contact the system's Primary User or Administrator to check on this.

The two most common errors messages are:

- `Host unreachable` - usually indicates a configuration error. It is caused by `arp` broadcasting the local network to convert an IP address to Ethernet address.
- `Network unreachable` - usually indicates no default or network route exists in the routing table.

Checking Additional Network Interfaces

If your system has more than one network interface (additional interfaces are usually fiber-optic [FDDI] links or SLIP connections or other Ethernet boards) you can easily perform the above checks on each network interface.

To check your other network interfaces, enter the `netstat -in` command. You see output similar to the following:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
ec0	1500	128.70.0	128.70.0.9	15	0	15	0	24
ec1	1500	128.70.2	128.70.2.5	15	0	15	0	24
sl0	1006	(pt-to-pt)	128.70.0.9	0	0	0	0	0
lo0	8304	loopback	localhost	8101	0	8101	0	0

The second Ethernet connection is to the network 128.70.2, a different LAN from the first Ethernet connection. The address of the local station on the second LAN is 128.70.2.5. To check that connection, use the `ping` command to test the connection to another station on that network.

There is also a SLIP link running in this example. The SLIP link extends the same LAN as `ec0` to another system in a different location. To test this link, find the hostname or IP address of the station at the other end of the SLIP link and use the `ping` command to test connectivity.

The `lo0` interface is the loopback network interface on the local host.

To determine which piece of hardware controls which interface, enter the command

```
ls -lS /hw/net/XXN
```

where `N` is the network interface number of type `XX`. For example:

```
% ls -lS /hw/net/ef0
crw----- 1 root sys 0 Jun 6 14:14 /hw/net/ef0 ->
/hw/module/4/slot/io1/baseio/pci/2/ef
```

Configuring a Dual-port Gigabit Ethernet Board

This section explains how to configure your system for Gigabit Ethernet. This configuration is recommended for a diskless workstation.

Note: The instructions in this section assume that your system already has (or has had) a functional Ethernet connection. If this is not the case, see the instructions in “Configuring a System for a Network” on page 39 on configuring your system for networking (giving it a hostname, IP address, and so on).

How To Change Your Primary Network Interface

To change your primary network interface, for example, to a Gigabit Ethernet board, follow these steps:

1. If necessary, become superuser (root), as follows:

```
% su
Password: thepassword
#
```

2. To determine your system's hostname, enter the following:

```
# hostname
```
3. Use your favorite editor to edit the `/etc/hosts` file. To open the file with `nedit`, enter

```
# /usr/sbin/nedit /etc/hosts
```
4. Find the line with your system's hostname and add the IP address entries. For example, if you are making an entry for a system with a hostname of `systemname` residing in the domain `group1.com` with a Gigabit Ethernet IP address of `187.8.27.6` and a Fast Ethernet IP address of `199.26.34.6`, enter these lines in the `/etc/hosts` file:

```
187.8.27.6 gate-systemname.group1.com gate-systemname #
gigabit ethernet

187.8.28.6 gate1-system.group1.com gate-systemname # gigabit
ethernet

199.26.34.6 systemname.group1.com systemname # fast ethernet
```
5. If your site uses an NIS service, make the changes in step 4 to the hosts database on the NIS server.
6. Open the file `/etc/config/netif.options` and find this line:

```
: if1name=
```
7. Change it to the following:

```
if1name=ef0 (tg0 on Origin 3000 series systems with i09)
```
8. Find this line:

```
: if1addr=
```
9. Change it to the following:

```
if1addr=systemname
```
10. Find this line:

```
: if2name=
```
11. Change it to the following:

```
if2name=tg0 tg1
```
12. Find this line:

```
: if2addr=
```

13. Change it to the following:

```
if2addr=gate-systemname
if3name = tg2
if2addr =gate1-systemname
```

14. Find this line:

```
: if_num=8
```

15. Change it to the following:

```
if_num=2
```

16. Save and close the file.

17. Restart the machine to ensure that the changes you just made will take effect.

Note: If you cannot shut down the machine, restart the network script.

Note: If you need to change the net mask or other options for the network interface cards, see “Subnetting a Network” on page 55

Introducing Network Management

This chapter provides brief overview information on managing a network after installation. Indications for further reading are followed by a summary of the network initiation and shutdown process. Introductions to a few basic network management tools are followed by pointers to help in interpreting network statistics. These specific topics are covered:

- “Resources for Network Management” on page 115
- “About Network Startup and Shutdown” on page 117
- “Network Management Tools” on page 119
- “Interpreting Network Statistics” on page 123
- “Troubleshooting Poor Network Performance” on page 129

Resources for Network Management

Network management is sufficiently complex an issue that a wide understanding of controls and “tuning knobs” is needed, far beyond any capacity to cover here. Suggested sources of background information include these books:

- *DNS and BIND*, by Paul Abitz and Cricket Liu, 2nd edition, O’Reilly and Associates.
- *Internetworking with TCP/IP Volume 1*, by Douglas E. Comer, 3rd edition, Prentice Hall, 1994.
- *LAN Management with SNMP and RMON*, by Gilbert Held, John Wiley and Sons, 1996.
- *Managing Internetworks with SNMP*, by Mark Miller, M & T Books, 1993.
- *Managing NFS and NIS*, by Hal Stern, 1st edition, O’Reilly and Associates, 1991.
- *Network Management - A Practical Perspective*, by Allan Leinwand and Karen Fang Conroy, 2nd edition, Addison Wesley, 1996.

- *Routing in the Internet*, by Christian Huitema, Prentice Hall, 1995.
- *System Performance Tuning*, by Mike Loukides, 1st edition, O'Reilly and Associates, 1990.
- *TCP/IP Network Administration*, by Craig Hunt, 1st edition, O'Reilly and Associates, 1992.
- *TCP/IP Illustrated, Volume 1*, by W. Richard Stevens, Addison Wesley, 1994.

Background reading should also include white papers and current periodicals on the subject because the field changes almost daily.

About Network Management Functions

Network management ranges from maintenance, monitoring, and controlling network events to problem isolation. Basic necessary components in network management tools and strategies are:

- Fault management—Replacing or repairing managed devices when they fail, keeping network error logs, and swapping failed components. A few methods of detecting failure are introduced in “Troubleshooting Poor Network Performance” on page 129.
- Accounting management—Identifying users of network resources and assigning departmental charges. This aspect of network management is covered in *IRIX Admin: Backup, Security, and Accounting*.
- Configuration and name management—Planning, initializing and updating network software and hardware. Previous chapters have dealt with some of these issues.
- Performance management—Ensuring network components perform to expectations, that the network uses available bandwidth appropriately, perhaps rewriting protocols and developing administration tools. Many aspects of these issues are covered in the pertinent application manuals. For example: *Performance Co-Pilot for ORACLE Administrator's Guide*, *Developer Magic: Performance Analyzer User's Guide*, and *IRIS HIPPI API Programmer's Guide*.
- Security management—Protecting managed objects from intrusion or misapplication. For security-related issues see *IRIX Admin: Backup, Security, and Accounting*.

Because many networks are hybrid in design and contain multiple protocols and a wide variety of devices, network management is the only cohesive ingredient in maintaining functionality and preventing chaos. The following sections on network startup and shutdown give a simple summary of how the network operates:

- “Network Initialization Process” on page 117
- “Network Shutdown Process” on page 119

About Network Startup and Shutdown

The main network script is `/etc/init.d/network`. Other scripts for other network applications (UUCP, mail, and so on) also reside in this directory, but are covered in their appropriate chapter in this guide. The `network` script is described briefly here.

The `network` master script is called during system startup and shutdown. It defines the system name and host ID, ensures that the system has a valid Internet address, starts networking daemons, and initializes the network interfaces. Previously, site-dependent configuration commands to start and stop local daemons, add static routes, and publish ARP entries were put in a separate shell script called `/etc/init.d/network.local`. Symbolic links were made from `/etc/rc0.d` and `/etc/rc2.d` to `/etc/init.d/network.local` so the `network.local` file could be called at system startup and shutdown (see “Creating a Local Network Script” on page 91 for setup procedure). Currently this function is performed by `/etc/config/static-route.options`.

The `network` master script is linked to `/etc/rc0.d/K40network`, which is invoked from `/etc/rc0` during shutdown, and to `/etc/rc2.d/S30network`, which is invoked from `/etc/rc2` during startup. The script understands two arguments: **start** and **stop**. It can be run manually for testing and troubleshooting network-related problems without having to reboot the system.

Network Initialization Process

During system initialization, the shell script `/etc/init.d/network` is called. These are the actions performed by the script at startup:

1. Checks hostname and Internet address to determine if system should be configured as standalone or networked. Checks *sys_id* and *hosts* files. If the **network** configuration flag is off, the system is configured for standalone operation.
2. Determines names and addresses of primary and router interfaces for typical configurations.
3. Obtains any site-dependent information for interfaces from the *netif.options* file.
4. If system is not diskless, the shell script flushes all old routes.
5. Configures all interfaces, including loopback, using the *ifconfig* command.
6. If configured for IP packet filtering, the shell script starts the IP packet filtering daemon (*/usr/etc/ipfilterd*). The *ipfilterd* daemon must be started before gateway interface initialization.
7. Initializes gateway interface.
8. Initializes additional interfaces specified in the *netif.options* file.
9. If specified, initializes the Hypernet interface according to the *ifconfig-hy.options* file.
10. Initializes the loopback interface.
11. Using the *chkconfig* command, determines daemon configuration and reads relevant daemon configuration files (**.options*).
12. Sets default route for all IP multicast packets to the primary interface.
13. If NIS software is configured, defines and sets NIS domain name.
14. If NIS software is configured, starts appropriate NIS daemons.
15. If NFS software is configured, starts appropriate NFS daemons and mounts any NFS filesystems listed in the */etc/fstab*.
16. If NFS software is configured, starts other NFS daemons if *autofs*, *cacheefs*, or *lockd* is configured on with *chkconfig*.
17. If configured on with *chkconfig*, it starts standard daemons (*inetd*, *timed*, *timeslave*, *rarpd*, *rwhod*, *snmpd*, and so on).

Network Shutdown Process

During system shutdown, */etc/init.d/network* stops the daemons and disables the network devices. These are the actions the script performs at system shutdown:

1. Kills all network services that may be associated with a shell (*rlogind*, *rexecd*, *rshd*, *ftpd*, *telnetd*, and so on).
2. Kills some network daemons immediately (*inetd*, *bootp*, *tftpd*, *snmpd*, and so on).
3. If NFS is running, unmounts remote filesystems.
4. Kills all remote daemons.
5. If NFS is running, unexports exported filesystems. See the *ONC3/NFS Administrator's Guide* and the *NIS Administration Guide* for complete information about the optional NFS software.
6. Kills daemons that must be kept alive until the last minute (*portmap*, *slip*, *ipfilterd*).
7. Gracefully takes the system off the FDDI ring, if it is on the ring.
8. Stops the *ypbind* process of NIS.

Network Management Tools

This section summarizes the most standard networking tools for day-to-day management of your network. Except as noted, these networking tools reside in the */usr/etc* directory. See the online reference pages for explicit instructions on how these tools are used. These standard networking tools are described briefly in the following subsections: *ifconfig*(1M), *netstat*(1), *arp*(1M), *rcpinfo*(1M), *ping*(1M), *spray*(1M), *rtquery*(1M), *traceroute*(1M), *route*(1M), *rup*(1C), *ttcp*(1), and *netsnoop*(1M). The optional SGI network management tool NetVisualyzer is also summarized.

ifconfig(1M) The *ifconfig* command sets or checks the network interfaces to TCP/IP; it assigns IP addresses, subnet mask, and broadcast address to each interface. *ifconfig* is performed at boot time by the master network configuration script */etc/init.d/network*. An example of modifying the *ifconfig-#.options* file is shown in Chapter 3. An example of using *ifconfig* to configure various interfaces is shown

in *TCP/IP Network Administration* by Hunt. Also see the `ifconfig(1M)` reference page for details on how to use the various options to this command.

`netstat(1)` The `netstat` command is used extensively to display which network interfaces are configured, which are available; `netstat` can also show whether a valid route to your destination is available. You can use `netstat` to display queue information (`-i q`), network memory (`-m`), and protocols (`-p`).

For sample output of `netstat` see Collecting Network Statistics with `netstat` on page 79. Examples of the use of `netstat` are given extensively in *TCP/IP Network Administration* by Hunt. An analysis of using `netstat` in the measurement of collisions and network saturation can be found in *Managing NFS and NIS* by Stern. Also see the `netstat(1)` reference page for details on how to use the various options to this command.

`arp(1M)` The `arp` command displays the contents of the ARP table which maintains dynamic translations between IP addresses and Ethernet addresses. This command can be used to detect systems on the local net that are configured with the incorrect IP address. `arp` options include `-a` for all entries, `-d` to delete an entry, `-s` to publish an entry and act as a server for this entry, and `-f` to pull information from a specified file instead of `/dev/kmem`. `arp`. The `arp(1M)` reference page describes the use of these options.

An example of using the `arp` command to troubleshoot incorrect address resolution is given in *TCP/IP Network Administration* by Hunt. `arp` does not display the local station's Ethernet address. To get a local station's Ethernet address, use the `netstat` command with the `-ia` options. A discussion of using `arp` to diagnose intermittent failures is given in *Managing NFS and NIS* by Stern.

`rpcinfo(1M)` The `rpcinfo` command queries Remote Procedure Call (RPC) servers and their registration with the portmapper; it is analogous to `ping` in that it verifies that the remote machine is capable of replying to an RPC request. `rpcinfo` can be used to detect and debug unresponsive servers, RPC client-server mismatches, and broadcast problems related to the RPC service. The information provided by `rpcinfo` includes a list of `rpc`-based applications (`portmapper`, `NIS`, `rstatd`, and so on), and the program number, version number, protocol (TCP/UDP), and associated port number. If you are running an `rpc`-based network

- application and cannot get a response from the remote station, use the *rpcinfo* tool to ensure that the remote station supports the desired application. A discussion of RPC mechanics can be found in *Managing NFS and NIS* by Stern. The *rpcinfo(1M)* reference page describes the use of available options.
- ping(1M)* The *ping* command tests whether a remote host is alive and reachable from your system. It is based on the Internet Control Message Protocol (ICMP) and sends an ECHO_REQUEST soliciting an ECHO_RESPONSE, thereby creating a two-way stream. It provides general information about packet loss and round-trip time.
- A discussion of using *ping* to test basic connectivity can be found in *TCP/IP Network Administration* by Hunt. A case study using *ping* to diagnose network health and an analysis of the results is given in *Managing NFS and NIS* by Stern.
- See also “Testing Network Connectivity With *ping*” on page 123 and the *ping(1M)* reference page.
- spray(1M)* The *spray* utility sends a one-way stream of packets to a station using consecutive packets of fixed length, making remote procedure calls to the *rpc.sprayd* daemon on the remote host to report how many packets were received, and the transfer rate. It roughly estimates network interface capacity both on individual hosts and through network hardware between hosts. Although it provides very limited information about general network performance, an example of using *spray* to gauge Ethernet interface capacity can be found in *Managing NFS and NIS* by Stern. See also the reference page for *spray(1M)*.
- rtquery(1M)* Sends a request to a designated station for information on the station’s network routing tables (*routed* or *gated*). This tool can be useful for troubleshooting routing problems; it shows the number of routing response packets returned. For details see RFC 1058—*Routing Information Protocol* and the *rtquery(1M)* reference page.
- traceroute(1M)* Track packets as they journey through the network with *traceroute*. This tool is very useful for isolating network and router faults in a large heterogeneous network. It displays the names and addresses of all the intermediary routers that support the Internet Protocol “time-to-live” (TTL) field. It also displays the amount of time the packet spends traveling to the router, on the router, and leaving the router. *traceroute*

- increases network load; this factor should be considered when testing a network with *traceroute*. See the *traceroute(1M)* reference page for an explanation of how this command uses options, and sample output.
- route(1M)** Use the *route* command to add or delete entries in the network routing tables. Typically, the routing tables are handled automatically by the *routed* or *gated* daemon. However, *route* can be used to create, maintain, and delete static routing tables, to flush routing tables, and to show metric information about routes. To have static routes incorporated at startup, modify the file */etc/gateways* and */etc/config/routed.options*.
- For an example of building a static routing table using the *route* command see *TCP/IP Network Administration* by Hunt.
- rup(1C)** Displays status information, including uptime and load average, about remote stations using Sun RPC broadcasts. If no specific station is specified, it uses broadcasting and returns information about stations on the local network; broadcasting does not go through routers. This tool is useful for isolating physical problems with a station or the network.
- ttcp(1)** Used to test Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) performance. This tool provides a more realistic measurement of performance than the standard tests (*spray*, *rup*, *ping*). It allows measurements to be taken at both the local and remote end of the transmission. See “Measuring Network Throughput With *ttcp*” on page 124.
- netsnoop(1M)** Captures packets and decodes network traffic, and in some ways acts as a network sniffer. You can either capture traffic to a file or watch real-time packets between any two machines. To capture or have a clear view of all traffic between two machines, *netsnoop* must be run on a third machine. In this way the view from the third machine will not be biased by what the sender thinks it is putting on the wire, but might not actually be getting there, nor by what the receiver is failing to receive from the wire.
- You can use *netsnoop* with a number of filters in realtime, or on a saved tracefile reactivated by the command *netsnoop*. You must be superuser to run *netsnoop* on a local network interface. *netsnoop* is a useful tool for analyzing damaged packets, and tracing problems. If network overload is indicated, NetVisualyzer (summarized below) includes tools such as *netcollect* that are designed to help in figuring out if subnetting is needed.

NetVisualyzer is network management tool available as an option for use on any SGI system. It is a passive network management product that offers a set of graphical traffic monitoring, diagnostics, planning, and performance analysis tools that provide network information and statistics for Ethernet or FDDI networks in a visually intuitive form. NetVisualyzer comprises six tools: NetLook, NetGraph, NetCPA, Analyzer, RouteQuery, and TraceRoute. NetVisualyzer allows you to view and monitor your network, collect network statistics and generate reports based on those statistics, and decode heterogeneous packets layer by layer.

Interpreting Network Statistics

The network management tools provide the network administrator with valuable information about the network. However, the presentation of these statistics can be overwhelming. This section illustrates how to use three of the most common management tools and how to interpret the network statistics generated by these tools:

- “Testing Network Connectivity With ping” on page 123
- “Measuring Network Throughput With ttcp” on page 124
- “Collecting Network Statistics With netstat” on page 126

Testing Network Connectivity With ping

Used with no options, the *ping* tool tells you if a machine is active and reachable over your network. You can also use *ping -s* to show datagrams of a given size and transit times (round trip times) and packet losses. An example of using *ping* to isolate network problems can be found in *Managing NFS and NIS* by Stern.

Unless used with the count *-c* option, *ping* continuously transmits datagrams until halted, and therefore should not be used in a script. When using *ping* to troubleshoot, start with the local host to ensure the local network interface is operational, then increase the radius of testing to remote hosts and gateways.

The following example uses *ping* with the *-c* option to limit the number of datagrams sent and received. The *ping* tool is testing and measuring traffic between the local station and the station *testcase*. See the ping(1M) reference page for more details about other *ping* options.

```
/usr/etc/ping -c5 testcase
PING testcase (192.55.43.4): 56 data bytes
64 bytes from 192.55.43.4: icmp_seq=0 ttl=249 time=160.314 ms
64 bytes from 192.55.43.4: icmp_seq=1 ttl=249 time=47.057 ms
64 bytes from 192.55.43.4: icmp_seq=2 ttl=249 time=28.129 ms
64 bytes from 192.55.43.4: icmp_seq=3 ttl=249 time=48.596 ms
64 bytes from 192.55.43.4: icmp_seq=4 ttl=249 time=131.894 ms
----testcase PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 28.129/83.198/160.314 ms
```

If *ping* shows no output, the pinged host is disconnected or out of service. Consistent packet loss over 0.1% should be investigated further. Damaged packets are sometimes an indication of faulty hardware, so tracing cabling, connectors, and terminators is a reasonable step. Packet loss may be sporadic, in which case a script to *ping* at specified intervals, saved to a file, would be useful. If *ping* responds but you still cannot *ftp* or *telnet*, you may be experiencing a TTL error. You can extend datagram time-to-live with the *-T* option. See the ping reference page for details.

Measuring Network Throughput With *ttcp*

The *ttcp* tool measures network throughput. It provides a realistic measurement of network performance between two stations because it allows measurements to be taken at both the local and remote ends of the transmission. As with all network management tools, the statistics must be interpreted with the network configuration and applications in mind. For example, the statistics generated from a *ttcp* probe between two stations with routers in between results in lower throughput than if the stations were located on the same network. On the same note, users running applications that transmit large data structures see slower throughput than users running applications that transmit smaller data structures.

You can use *ttcp* by issuing the following commands, the first from the receiver

```
ttcp -r -s -l 32768
```

the second from the sender:

```
ttcp -t -s -l 32768 -n XXX host
```


The receiver should be started before the sender. The `-l` option sets the length of the write/read size to 32K chunks. You can change this to 64K or 128K for slightly better performance, as long as it is a multiple of the page size. `ttcp` does memory-to-memory TCP data transfer between end systems, and measures end-to-end TCP throughput.

You should get over 800 KB per second on a 10baseT half duplex network and over 1000 Kbytes per second on a 10baseT full duplex switched network. On a 100baseT network you should get over 8500 KB per second for half duplex, and over 11500 KB per second for full duplex.

Most of the speed difference between half and full duplex is related to the Ethernet capture effect rather than anything else. There should be no `netstat -i` input nor output errors. Some of the network device drivers support a `ifconfig xxx debug` mode which causes more informative diagnostic error messages to be displayed on the console when errors occur. Use `ifconfig xxx -debug` to disable them again.

The following example illustrates the statistics you might see if you ran a simple `ttcp` test between the stations *sheridan* and *longstreet* (two workstations) on a clean network. See the `ttcp(1)` reference page for details about the many `ttcp` options.

On *sheridan*, enter the command

```
ttcp -r -s
```

You see the following output:

```
ttcp-r: buflen=8192, nbuf=2048, align=16384/0, port=5001 tcp
ttcp-r: socket ttcp-r: accept from 192.102.108.4
ttcp-r: 16777216 bytes in 19.99 real seconds = 819.64 KB/sec +++
ttcp-r: 10288 I/O calls, msec/call = 1.99, calls/sec = 514.67
ttcp-r: 0.1user 3.4sys 0:19real 17%
```

On *longstreet*, enter the command

```
ttcp -t -s sheridan
```

You see the following output:

```
ttcp-t: buflen=8192, nbuf=2048, align=16384/0, port=5001 tcp -> sheridan
ttcp-t: socket
ttcp-t: connect
ttcp-t: 16777216 bytes in 19.98 real seconds = 820.02 KB/sec +++
ttcp-t: 2048 I/O calls, msec/call = 9.99, calls/sec = 102.50
ttcp-t: 0.0user 2.3sys 0:19real 12%
```

The throughput statistics are highlighted in bold and are in units of KBps. The throughput on the station *sheridan* is 819.64 KBps and the throughput on the station *longstreet* is 820.02 KBps. Both throughput values indicate good network performance between the stations.

Collecting Network Statistics With netstat

The *netstat* tool shows the status of the network by identifying active sockets, routing and traffic. You can see network addresses for the interfaces, and the maximum transmission unit (MTU) depending on the option selected. On a given system, if the output collision rate is consistently greater than 5% of that on other systems, it may indicate a suspicious physical fault—a bad tap, transceiver, loose terminator, and so forth. Network-wide collision rate is discussed later.

An analysis of using *netstat* in the measurement of collisions and network saturation can be found in *Managing NFS and NIS* by Stern. The example given here illustrates the statistics you might see on a station using *netstat*. See the `netstat(1)` reference page for details about the many *netstat* options:

```
netstat -i
```

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
enp0	1500	central	thumper	498690	937	1066135	3	4858
lo0	32880	loopback	localhost	1678915	0	1678915	0	0

The collision rate is approximately 0.45%, well within the acceptable range.

If the network interface is assigned IP aliases, the output from `netstat -i` looks similar to this:

netstat -i

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
enp0	1500	central	thumper	498690	937	1066135	3	4858
enp0	1500	bldg-2	hopper	584280	163	137245	0	40717
lo0	32880	loopback	localhost	1678915	0	1678915	0	0

If the network-wide collision rate is consistently above 10%, it may be time to consider subdividing the network, although there are other possible culprits in the scenario—insufficient memory on diskless workstations, among several. See “Troubleshooting Network Daemons to Improve Network Performance” on page 131. Throughput is a reliable indicator of network performance (see “Measuring Network Throughput With `ttcp`” on page 124).

Network Tuning Information

In general, IRIX TCP/IP services are factory-tuned for optimum performance under routine network operating conditions and no tuning is required after a network setup. Kernel TCP/IP tuning parameters are controlled with the `sysune` command or from the file `/var/sysgen/master.d/bsd`.

Caution: It is not advisable to change the factory-shipped settings in this file unless you receive explicit instructions in product documentation or you are highly experienced in network administration.

The following subsections deal with specific tuning aspects of your network:

- “About Setting MTU Sizes” on page 128
- “About Setting Packet Forwarding” on page 128
- “About Setting Window Sizes” on page 128
- “HTTP Considerations” on page 129

About Setting MTU Sizes

IRIX 6.2 and later versions implement MTU Discovery, a feature that allows a host to calculate maximum transmission unit (MTU) sizes using a table of commonly used MTU sizes. MTU Discovery is specified by the *tcp_mtudisc* variable in */var/sysgen/master.d/bsd*. When *tcp_mtudisc* is on and a router requests packet fragmentation without providing MTU size information, the calculated segment size is used to fragment the packet instead of the default maximum segment size (specified by *tcp_mssdflt*). The MTU Discovery flag is on (set to 1) by default.

About Setting Packet Forwarding

By default, IRIX hosts that contain more than one network interface are assumed to be routers, so IP packet forwarding is enabled at system startup. In addition, the factory-shipped configuration enables ICMP redirection, which causes the router to forward messages that are destined for a particular host to an alternative router instead. Two tunable kernel parameters, *ipforwarding* and *icmp_droptoredirects*, turn packet forwarding and ICMP redirection on and off.

On firewall systems, *ipforwarding* and *icmp_droptoredirects* must be turned off to prevent packets from breaching the security boundary. See the product documentation for your firewall application for specific instructions on setting IP forwarding characteristics. Also see the *sysstune(1M)* reference page for more information.

About Setting Window Sizes

Send and receive window sizes are specified by the variables *tcp_sendspace* and *tcp_recvspace* in */var/sysgen/master.d/bsd*. By default, both variables are set to 60 KB. For wide area network traffic, reducing window sizes to a value in the range of 4 to 8 KB can result in better response times in interactive sessions. Large amounts of bulk transfer traffic from systems using TCP windows larger than 4 to 8 KB can overflow the buffers of slow- or moderate-speed routers on wide area network. Overflowing WAN router buffers reduces the speed of the bulk transfers, and interactive traffic usually must wait behind such large bursts in a route queue.

Note: Using small window sizes on high-speed WAN links (such as T3) severely reduces TCP speed.

On a local Ethernet network, reducing the window size often increases performance. However, it degrades performance on an FDDI network.

HTTP Considerations

In some cases, variables in `/var/sysgen/master.d/bsd` need adjusting for HyperText Transport Protocol (HTTP) servers. These variables frequently include changes to the number of connections and connection time-outs. See the product documentation for your HTTP application for specific tuning instructions.

Troubleshooting Poor Network Performance

Given the complexity of most networks it is often hard to troubleshoot poor network performance. Among the many factors to suspect are hardware problems, outgrown network configuration, heavy use of particular network applications, and excessive packet size. In general `/usr/adm/SYSLOG` is a good place to start looking for diagnostic messages. The following subsections consider particular aspects of network performance improvement:

- “Troubleshooting Hardware Problems to Improve Network Performance” on page 129
- “Troubleshooting Network Configuration to Improve Network Performance” on page 130
- “Troubleshooting Network Daemons to Improve Network Performance” on page 131
- “Decreasing Packet Size to Improve Network Performance” on page 131
- “Kernel Configuration for Better Network Performance” on page 132

Troubleshooting Hardware Problems to Improve Network Performance

A network can be slow or inoperable due to hardware malfunction, often manifesting in the form of packet loss or corruption. This can increase network traffic to the point of unmanageable congestion. Start by segmenting the network, and begin diagnosis section by section. The most obvious items to check at the physical level are these:

Controller board

Even if the network media bandwidth is capable of handling the network traffic load, the individual station may not be able to handle the traffic. This is evidenced by a high degree of traffic on the network interface for no apparent reason. This traffic can be seen using the *gr_osview* tool (see the *gr_osview(1)* online reference page for options to see network traffic statistics). If traffic is unusually heavy on the interface, then there may be a problem with the controller, or the controller may be too slow to handle the volume of traffic. You may need a high-speed controller like the Efast card.

Transmitter and controller

Ensure that the Signal Quality Error (SQE), also called *heartbeat*, is disabled on both the transmitter and controller. SQE can cause unnecessary network traffic between the local station and the transceiver. See the installation guides for your network controller and transceiver for instructions on disabling SQE. By default, all SGI network controller boards are shipped with SQE disabled.

Physical problems with the media

Cables, taps, and other hardware will periodically break or malfunction. A time domain reflectometer (TDR) is essential for troubleshooting Ethernet cable problems. A good analyzer is also strongly recommended to assist in isolating network physical problems. SGI NetVisualizer product supplies a visual network analyzer ideal for locating physical problems with the media. Disk access difficulties can also contribute to slow network performance. An otherwise healthy network feels sluggish if NFS requests requiring disk access are not handled promptly.

Troubleshooting Network Configuration to Improve Network Performance

The network configuration or topology may also be a factor in adverse network performance. Monitor network usage to benchmark when the network is performing well. It provides a standard against which any changes can be measured. In summary check:

- Physical path—the wiring (if appropriate) for EMI and crosstalk.
- Concentration points—hubs and concentrators for capacity and connectivity.
- Number and location of repeaters and bridges —their age and intelligence.

- Router and gateway configurations— for traffic load and throughput, particularly if connecting large subnets ensure that there are at least two ways in and out of a network.
- Work group affiliations and routine sharing of resources (NFS filesystems, NIS domain, electronic mail, financial databases).
- Client/server configurations and the traffic they generate.
- Individual workstation usage, specific disk activity, burden of directory name lookups.

The complexity of network configuration requires a familiarity with all of the resources listed in “Resources for Network Management” on page 115.

Troubleshooting Network Daemons to Improve Network Performance

Some network daemons (*rwhod*, *rtnetd*, and so on) can have an undesirable effect on the network or network interface. For example, if a workstation is a multi-processor or is running real-time processes, the *rtnetd* daemon may be running on the station. This daemon is responsible for preempting incoming network packets to provide better response time for real-time processes. This is perfectly acceptable if the user is aware of the trade-offs between network processing and real-time processing. Some network daemons should be evaluated individually.

Do not load *rtnetd* software on routers or other network intensive stations (mail servers, NIS and DNS servers, and so on).

Decreasing Packet Size to Improve Network Performance

The maximum transfer unit (MTU) for data on the Ethernet is 1500 bytes. Network performance and efficiency increase with packet size up to the MTU for the medium. Packets that are larger than the media’s MTU must be broken into smaller packets (fragmented) to fit within the medium’s MTU. IRIX versions 6.2 and higher use MTU Discovery to perform fragmentation (see “About Setting MTU Sizes” on page 128 for details).

Kernel Configuration for Better Network Performance

You can change several parameters to customize network behavior for local configurations. The parameters listed below are in the `/var/sysgen/master.d/bsd` configuration file. For details on reconfiguring the kernel after changing this file, see *IRIX Admin: System Configuration and Operation*.

Changing Kernel Tunable Options to Improve Network Performance

There are four kernel parameters that directly affect the network performance: `tcp_sendspace`, `tcp_recvspace`, `udp_sendspace`, and `udp_recvgrams`.

These parameters determine the default amount of buffer space used by TCP (SOCK_STREAM) and UDP (SOCK_DGRAM) sockets. The `tcp_sendspace` and `tcp_recvspace` parameters define the initial buffer space allocated to a socket. The `udp_sendspace` parameter defines the default maximum size of UDP datagrams that can be sent. The `udp_recvgrams` parameter determines the number of maximally sized UDP datagrams that can be buffered in a UDP socket. The total receive buffer size in bytes for each UDP socket is the product of `udp_sendspace` and `udp_recvgrams`. A program can increase or decrease the send buffer and receive buffer sizes for a socket with the `SO_SNDBUF` and `SO_RCVBUF` options to the `setsockopt` system call. Many older TCP implementations have problems with large TCP `sendspace/recvspace` values. This should be decreased from 60 to 24 in environments where older stations have problems communicating.

For 4.2BSD compatibility, the IRIX system limits its initial TCP sequence numbers to positive numbers.

Changing TCP Parameters to Improve PC Connectivity

Many industry-standard personal computers with TCP/IP implementations experience difficulty connecting to SGI workstations and servers. This is because of the increased size of the `tcp_sendspace` and `tcp_recvspace` variables in the IRIX file `/var/sysgen/master.d/bsd`.

To allow your personal computers to connect successfully, change the values of the above variables from the default (60 * 1024) to (24 * 1024) and reconfigure the kernel with the `lboot` command. For more information on reconfiguring these values, see *IRIX Admin: System Configuration and Operation*.

SLIP and PPP

This chapter introduces the SGI implementation of the Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP). SLIP and PPP are protocols for TCP/IP networking over a serial line. SLIP and PPP can be used for connecting remote systems to a local area network, or for connecting two networks together.

The following sections are included in this chapter:

- “About SLIP and PPP” on page 134
- “Verifying the SLIP and PPP Software” on page 136
- “Selecting a Modem” on page 137
- “IP Addresses for SLIP and PPP Clients” on page 137
- “Configuring a System for Dial-Out” on page 138
- “Configuring a System for Dial-In” on page 144
- “SLIP and PPP Routing and Address Allocation” on page 146
- “Configuring a Bidirectional Link” on page 150
- “Starting SLIP or PPP at Boot Time” on page 150
- “About Demand Dialing” on page 151
- “NFS Over SLIP or PPP” on page 152
- “File Transfer Over SLIP or PPP” on page 152
- “Troubleshooting SLIP and PPP Links” on page 152

About SLIP and PPP

The SGI implementation of SLIP and PPP provides both RFC 1144 data compression and its own proprietary data compression, which compresses header framing, checksum, and TCP/IP information to three bytes. SLIP is part of the *eo*e software subsystem, and is installed via the *eo*e.*sw*.*slip* package. Use the *versions* command to see if you have SLIP installed on your system.

PPP offers substantially the same features as SLIP, but it is more flexible and robust. For this reason, PPP is usually preferred over SLIP (however, you must use SLIP if the system to which you are connecting supports SLIP and not PPP.) In addition to its use in dial-up connections over telephone lines and modems, PPP is also used in ISDN connections; however, this guide assumes that you are using PPP over a conventional serial line connection.

Note: If you are connecting a remote host with ISDN, you might prefer to use the configuration instructions in the *ISDN User's Guide* instead of using the procedures in this chapter (also see the *isdn(1M)* reference page). You should use the information in this chapter for an understanding of PPP operation and maintenance, however.

PPP is part of the *eo*e software subsystem, and is installed via the *eo*e.*sw*.*ppp* package. Use the *versions* command to check to see if you have PPP installed on your system.

Usually, a SLIP or PPP connection serves one of the following purposes:

- To connect a single remote system to a network. Here the systems have a client-server relationship, with the single system being the client, and the system directly connected to the network being the server. The server handles all of the routing.
- To connect two networks together. In this case, the two systems usually act as peers.

A given SLIP or PPP link can be set up so that one system always initiates the connection (by dialing the other system), or so that either system can initiate the connection. In the case of a client-server connection, the client usually initiates the connection by dialing the server. However, the connection can be bi-directional as well. In the case of a connection between two networks, you can either choose one system to initiate connection, or allow either system to initiate the connection as needed.

Setting Up a SLIP or PPP Connection: General Procedures

Setting up an individual system as a SLIP or PPP client requires the following steps:

1. Check that you have the SLIP and PPP software installed. See “Verifying the SLIP and PPP Software” on page 136 to find out whether you have the software installed.
2. Install the software if necessary. See “Installing the SLIP and PPP Software” on page 137 if you need to install software.
3. Select and install the appropriate hardware (modem and cable). See “Selecting a Modem” on page 137 for further information.
4. Select an IP address. IP addressing alternatives are described in “IP Addresses for SLIP and PPP Clients” on page 137.
5. Configure the software for dial-out. See “Configuring a System for Dial-Out” on page 138.
6. Configure routing behavior. Various routing alternatives are described in “SLIP and PPP Routing and Address Allocation” on page 146.

Setting up a system as a SLIP or PPP server requires the following steps:

1. Check that you have the SLIP and PPP software installed. See “Verifying the SLIP and PPP Software” on page 136.
2. Install the software, if necessary. See “Installing the SLIP and PPP Software” on page 137.
3. Select and install the appropriate hardware (modem and cable). See “Selecting a Modem” on page 137 for further information.
4. Configure the software for dial-in (set up configuration information for each client). See “Configuring a System for Dial-In” on page 144.
5. Configure routing behavior. Various routing alternatives are described in “SLIP and PPP Routing and Address Allocation” on page 146.

Connecting two networks using SLIP or PPP requires the following steps:

1. Check that you have the SLIP and PPP software installed. See “Verifying the SLIP and PPP Software” on page 136.
2. Install the software on both systems if necessary. See “Installing the SLIP and PPP Software” on page 137.

3. Select and install the appropriate hardware (modem and cable) on both systems. See “Selecting a Modem” on page 137.
4. Configure the software for dial-out access on one or both systems. See “Configuring a System for Dial-Out” on page 138.
5. Configure the software for dial-in access on one or both systems. See “Configuring a System for Dial-In” on page 144.
6. Configure routing behavior. Various routing alternatives are described in “SLIP and PPP Routing and Address Allocation” on page 146.

Verifying the SLIP and PPP Software

Make sure you have the appropriate software installed: to use SLIP, you must have the *oe.sw.slip* and *oe.sw.uucp* subsystems installed. To use PPP, you must have the *oe.sw.ppp* and *oe.sw.uucp* subsystems installed. You should also install the corresponding reference pages (*oe.man* subsystems). To verify whether these packages are installed, use the *versions* command. For example, to check whether you have PPP installed, use the following command:

```
% versions ooe.\*.{ppp,uucp}
```

If PPP and UUCP are installed, *versions* should produce output something like this:

I = Installed, R = Removed

	Name	Date	Description
I	oeo	09/24/96	IRIX Execution Environment, 6.3
I	oeo.man	09/24/96	IRIX Execution Environment Man Pages
I	oeo.man.ppp	09/24/96	Point-to-Point Protocol Man Pages
I	oeo.man.uucp	09/24/96	UNIX-to-UNIX Copy Man Pages
I	oeo.sw	09/24/96	IRIX Execution Environment Software
I	oeo.sw.ppp	09/24/96	Point-to-Point Protocol Software
I	oeo.sw.uucp	09/24/96	UUCP Utilities

Installing the SLIP and PPP Software

For instructions on installing software, see *IRIX Admin: Software Installation and Licensing*. After installing the SLIP or PPP software, you must reboot your system to reconfigure the kernel.

Selecting a Modem

A modem capable of at least 9600 bits per second (bps, sometimes referred to as “baud”) is required for use with SLIP or PPP. However, at this speed, your link will be quite slow. A good modem choice is any modem that supports the V.32bis or V.34 standards and a speed of 14,400 bps or greater. A half-duplex modem can be used but is less desirable for interactive tasks.

SGI does not manufacture these modems and cannot be responsible for changes to the modems. SGI cannot guarantee modem compatibility or the quality of the telephone line used with SLIP or PPP, but a separate data-grade telephone line is highly recommended.

If you are using a high-speed modem, be sure it supports hardware flow control using RTS/CTS. With such a modem, use the `tttyf*` device name.

Note: SGI supports DSI, Inte, Telebit, ZyXEL U.S. Robotics, Hayes, and most Hayes compatible modems. While many others work, configuring them may be much more complicated and their operation is not guaranteed.

For instructions on installing modems and configuring modems, see “Installing a Modem” in *IRIX Admin: Peripheral Devices*.

IP Addresses for SLIP and PPP Clients

Each SLIP or PPP client needs an IP address and hostname for its SLIP or PPP interface. If it is also connected to one or more local area networks, it will also require an address and hostname for each of its other network interfaces. If you have had an address and hostname assigned to you by your system administrator or network service provider, you can use this information to configure your system.

Some network service providers use dynamic addressing, where each client is given a dynamically allocated IP address when it connects. IRIX PPP has support for dynamic addressing, but IRIX SLIP does not support dynamic addressing. See “Using Dynamic Address Allocation With PPP” on page 150 for more information about dynamic addressing.

If you are setting up a SLIP or PPP server on your LAN, or are connecting two networks over a serial link, you should see “SLIP and PPP Routing and Address Allocation” on page 146 for planning information.

Configuring a System for Dial-Out

Configuring SLIP or PPP for dial-out requires approximately the same steps:

1. Add a line to the */etc/uucp/Systems* file describing the connection.
2. Add a line for the modem in the */etc/uucp/Devices* file.
3. If your modem is not already listed in the */etc/uucp/Dialers* file, add an entry for it.
4. Verify that the modem line is correctly configured in the */etc/inittab* file.

PPP requires one additional step, setting up the */etc/ppp.conf* file.

These sections explain how to configure SLIP and PPP for dial-out:

- “Configuration Files for Dial-Out” on page 138
- “Sample SLIP Configuration for Dial-Out” on page 142
- “Sample PPP Configuration for Dial-Out” on page 142

Configuration Files for Dial-Out

This section describes the various configuration files for SLIP and PPP:

- */etc/uucp/Systems* is described in “Configuring */etc/uucp/Systems* for Dial-Out” on page 139.
- */etc/uucp/Devices* is described in “Configuring */etc/uucp/Devices* for Dial-Out” on page 140.

- `/etc/uucp/Dialers` is described in “Configuring `/etc/uucp/Dialers` for Dial-Out” on page 140.
- `/etc/uucp/inittab` is described in “Configuring `/etc/inittab` for Dial-Out” on page 141.
- `/etc/ppp.conf` is described in “Configuring `/etc/ppp.conf` for Dial-Out” on page 141.

Configuring `/etc/uucp/Systems` for Dial-Out

The `/etc/uucp/Systems` file contains the information your system needs to dial up another system. The remote station’s node name and telephone number, as well as the local modem’s speed and a password, are all kept here and are used to log in to the remote station. The format for a line representing a SLIP or PPP connection in the systems file is

```
system Any type speed phone login-script
```

`system` specifies the name of the remote system. `type` gives the name of an entry in the `Devices` file, specifying a line and modem type. `speed` specifies the speed of the connection between the system and the modem. `phone` specifies the phone number for the remote system. The `login-script` tells SLIP or PPP how to log into the remote system.

In this example, the local station “wenders” uses this line in its `/etc/uucp/Systems` file to call station lynch. The connection is made at 38,400 bps. The password is “hopper.” SLIP logs in to the remote station by responding with “slip-wenders” to the login prompt, and “hopper” to the password prompt.

```
lynch Any ACUSLIP 38400 5551212 "" \r\c ogin:--ogin: slip-wenders \
asswd: hopper SLIP
```

The information must be in one continuous line. The last string, `SLIP`, forces the local station to wait for the remote station to announce that it is starting the SLIP protocol.

The third field in the systems file (in this case, `ACUSLIP`) specifies the modem line to be used to call the remote station. There must be at least one entry matching this field in the `/etc/uucp/Devices` file.

For more information on the `/etc/uucp/Systems` file, see “UUCP Systems File” on page 219.

Configuring `/etc/uucp/Devices` for Dial-Out

The file `/etc/uucp/Devices` is used to configure the desired device, modem speed, and dialer program for SLIP on your IRIS station. The correct format for a line appropriate for SLIP in `/etc/uucp/Devices` is

```
type device null speed 212 x dialer
```

The first field, *type*, can be any string you like, but it should correspond with the type specified in the Systems file. If you have multiple modems that can be used interchangeably, they should all have the same *type* name—the system automatically selects one of these modems when making a call. The *type* name ACUSLIP is commonly used to designate a modem for use with SLIP or PPP.

device can be any flow-control device associated with a port not currently in use. With high-speed modems, hardware flow control and a sound cable is recommended. *speed* should be the speed of the connection between your system and the modem. *dialer* can be any dial program listed in `/etc/uucp/Dialers`.

If you want to configure more than one modem speed, use a different port, or use a modem that supports a different command set, create a new line in the devices file that reflects the change.

For example, the following line configures SLIP to use a Telebit™ T2500 modem at 38,400 bps on the serial port 2 hardware flow control device:

```
ACUSLIP ttyf2 null 38400 212 x t25slip
```

For more information on the `/etc/uucp/Devices` file, see “UUCP Devices File” on page 211.

Configuring `/etc/uucp/Dialers` for Dial-Out

The `/etc/uucp/Dialers` file contains entries for various types of modems. The *dialer* field in the `Devices` file should refer to one of these entries.

It should not be necessary to add an entry to the dialers file, unless you are installing a type of modem not supported by IRIX. If you need to add or modify an entry in the dialers file, see “UUCP Dialers File” on page 216.

Configuring `/etc/inittab` for Dial-Out

The port you specified in `/etc/uucp/Devices` must be configured for dial-out or dial-in/dial-out use in the `/etc/inittab` file. If you set up your modem for dial-out or dial-in/dial-out use as directed in “Installing a Modem” in *IRIX Admin: Peripheral Devices*, the `/etc/inittab` file should already be set up correctly.

For example, if you want to use `ttyf2` for dial-out SLIP, the line for `ttyf2` in `/etc/inittab` should read

```
t2:23:off:/etc/getty ttyf2 co_38400          # port 2
```

This line turns off the `getty` program on port number 2.

If the link can be initiated by either station, you must turn on `uugetty`. For example, to configure a symmetric link using Telebit T2500 modems, change the line to

```
t2:23:respawn:/usr/lib/uucp/uugetty -Nt 60 -it25in,conn ttyf2 dx_38400
```

Changes made to `/etc/inittab` are acted upon when `init` reexamines the `/etc/inittab` file. To make `init` reexamine `/etc/inittab` immediately, use this command:

```
/etc/telinit q
```

For more information, see `inittab(4)`.

Configuring `/etc/ppp.conf` for Dial-Out

The `/etc/ppp.conf` file is used to specify options for PPP connections. Each entry in the file consists of a host name and a set of options. For example:

```
salad      out remotehost=dial-in.salad.com
           localhost=caesar.salad.com
           quiet add_route
```

This example specifies a connection between the local host `caesar.salad.com` and a remote machine called `dial-in.salad.com`. The **out** keyword specifies that this is an outgoing connection. The **quiet** keyword specifies a demand-dialed connection (see “About Demand Dialing” on page 151). The **add-route** keyword tells PPP to set up a default route through `dial-in.salad.com`.

For a complete list of options for the `ppp.conf` file, see the `ppp(1M)` reference page.

Sample SLIP Configuration for Dial-Out

This section shows a sample SLIP configuration for dial-out. This example sets up a connection between tuna.salad.com and dial-in.salad.com, a server on the salad.com corporate network.

```
/etc/uucp/Systems
salad Any ACUSLIP 38400 5551212 "" \r\c ogin:--ogin: slip-tuna \
asswd: celery SLIP
```

```
/etc/uucp/Devices
ACUSLIP ttyf2 null 38400 212 x t25slip
```

```
/etc/inittab
t2:23:off:/etc/getty ttyf2 co_38400          # port 2
```

To start SLIP using the configuration in this example, enter the following command:

```
% /usr/etc/slip -o -p comp -r salad
```

Sample PPP Configuration for Dial-Out

The main configuration file for PPP is */etc/ppp.conf*. This file is described in detail in the ppp(1M) reference page. Here is an example of a *ppp.conf* file:

```
salad      out remotehost=dial-in.salad.com
           localhost=caesar.salad.com
           quiet add_route
```

This entry describes an outgoing connection from a standalone system to a remote host named “dial-in.salad.com,” which acts as a gateway to the spice.com network. The entry specifies demand-dialed (**quiet**) mode. Stand-alone clients should usually include the **add_route** keyword to set up a default route through the PPP server. The entry *salad* should match an entry in the */etc/uucp/Systems* file, such as the following:

```
salad Any ACUSLIP 38400 555-1212 "" @\r\c ogin:--ogin: ppp-caesar \
ssword: mypasswd PPP
```

IRIX PPP sends out the message `starting PPP` before starting the protocol; therefore, the chat script shown above waits to receive the string PPP to ensure that the login has succeeded. If you are connecting to a non IRIX system, you may need to remove the PPP string from this example.

There must be at least one ACUSLIP entry in the `/etc/uucp/Devices` file:

```
ACUSLIP ttyf2 null 38400 212 x t25slip
```

This entry specifies that `ttyf2` is connected to a Telebit 2500 modem.

You should also have your `/etc/inittab` file set up to expect the modem port speed you are using (in this case, 38400 bps) and to have `getty` or `uucpgetty` turned off. Specific details on editing the `/etc/inittab` file and restarting `telinit` are found in *IRIX Admin: Peripheral Devices*. The following entry works with the above listed file entries for PPP:

```
t2:23:off:/etc/uucp/uucpgetty ttyd2 dx_38400 # ppp modem
```

You may use dial-out PPP with the above-listed entry, but for dial-in PPP, you must configure `uucpgetty` to answer the line, as described in *IRIX Admin: Peripheral Devices*.

It is important to note that the above-listed entries are simply examples of one configuration that works for one example site. The same entries may not yield satisfactory results in every case, due to other differences in site configuration and modem manufacturer and model. For example, the PPP site you are dialing into may require different settings in the `/etc/ppp.conf` file, and the entry in the `Devices` file assumes a specific brand and model of modem. Also, the example assumes that you have configured the modem as described in "Installing a Modem" in *IRIX Admin: Peripheral Devices*.

When you have configured the files, you must enter the `ppp` command as `root`. The following command works for the above-listed example file entries:

```
ppp -r salad
```

For complete information on the `ppp` command and its options, see the `ppp(1M)` reference page.

Configuring a System for Dial-In

To configure any system for dial-in, you must have at least one port set up for dial-in use (or combination dial-in/dial-out use). Each system that can dial in requires an entry in the `/etc/passwd` file. In addition, SLIP connections may require an entry in the `/usr/etc/remoteslip` file, and PPP connections may require an entry in `/etc/ppp.conf`.

These sections describe how to configure SLIP and PPP for dial-in:

- “Configuring `/etc/passwd` for Dial-In With SLIP” on page 144
- “Configuring `/usr/etc/remoteslip` for Dial-In With SLIP” on page 145
- “Configuring `/etc/passwd` for Dial-In With PPP” on page 146

Configuring `/etc/passwd` for Dial-In With SLIP

SLIP requires an entry in `/etc/passwd` in order to log in. The user ID and group ID must both be zero (0). Instead of the shell specified at the end of a normal entry in `/etc/passwd`, SLIP uses the file `/usr/etc/remoteslip`. To allow `tuna.salad.com` to log in as “slip-tuna,” `dial-in.salad.com` should have this line in `/etc/passwd`:

```
slip-tuna:3RsB768WRAN2.:0:0:slip for tuna:/:usr/etc/remoteslip
```

An entry like this one is necessary for each station that calls in by using SLIP. For maximum system security, the home directory for SLIP accounts should only be writable by the superuser. Using open directories such as `/tmp` opens your system up to a variety of threats.

Note: The encrypted password in the example does not represent a real password. You must use the `passwd` command to set the password for the SLIP login. See the `passwd(1)` reference page for information on using `passwd`.

Configuring `/usr/etc/remoteslip` for Dial-In With SLIP

In the `/etc/passwd` entry for `slip-tuna`, the login shell is specified as the file `/usr/etc/remoteslip`. This file is used to invoke SLIP on a remote station. In `/usr/etc/remoteslip`, the `slip` command can specify the remote station's name and any other options appropriate to that connection.

`/usr/etc/remoteslip` is a Bourne shell script. You can add to the case statement as you would to any Bourne shell script case statement. The `sh(1)` reference page contains detailed information about shell script programming. Each SLIP connection should have an entry in the following format:

```
slip-nodename )
    exec /usr/etc/slip options
    ;;
```

`nodename` is the name of the remote station. Options for `slip` are detailed in the `slip(1M)` reference page.

The `/usr/etc/remoteslip` file might contain this entry on the station `dial-in.salad.com`:

```
# Edit the case statement as required.
case $USER in
    slip-tuna)
        exec /usr/etc/slip -p comp -i -r tuna
        ;;
    *)
        exec /usr/etc/slip -i -r $USER
        ;;
esac
```

The connection between `dial-in.salad.com` and `tuna` uses the SGI proprietary header prediction and compression for faster data transfer because the `-p comp` option is specified. To use RFC 1144 compression, use `-p cslip` instead. The option tells SLIP that the session is input from another station. The `slip` option, `-r tuna`, specifies the remote station's name. Note that this example also supplies a default case. If all of your SLIP clients use the same parameters, you can just modify the default case, instead of adding entries for each client.

Configuring `/etc/passwd` for Dial-In With PPP

Like SLIP, PPP requires an entry in `/etc/passwd` in order to log in. The user ID and group ID must both be zero (0). Instead of the shell specified at the end of a normal entry in `/etc/passwd`, PPP uses the command `/usr/etc/ppp`. To allow the system `caesar.salad.com` to log in as “`ppp-caesar`,” `dial-in.salad.com` should have this line in `/etc/passwd`:

```
ppp-caesar:3RsB768WRAN2.:0:0:PPP for caesar:/:usr/etc/ppp
```

An entry like this one is necessary for each station that calls in by using PPP. For maximum system security, the home directory for PPP accounts should be writable only by the superuser. Using open directories such as `/tmp` exposes your system to a variety of threats.

Note: The encrypted password in the example does not represent a real password. You must use `passwd` to set the password for the PPP login. See the `passwd(1)` reference page for information on using `passwd`.

If the default parameters are acceptable, you don't even need an entry in `/etc/ppp.conf` for a client. You may want to add a minimal entry, such as this:

```
ppp-caesar      in remotehost=caesar.salad.com
```

Note: You should never use the **add-route** keyword by itself on a server system.

SLIP and PPP Routing and Address Allocation

There are three basic ways to deal with routing over a SLIP or PPP link, depending on the circumstances.

- If you have a few standalone clients connecting to a server with SLIP on your main network, you probably want to assign client addresses from the server's network, and use proxy-ARP routing (PPP handles ARP table entries automatically). To learn more about proxy-ARP routing, see “About Proxy-ARP Routing for SLIP Connections” on page 147.

- If you have a lot of standalone clients connecting to your server, you're better off setting aside a subnet for SLIP and PPP client addresses. In this case, the server appears to the main network as a gateway to the SLIP/PPP "net." This issue is explored in "Setting Up SLIP/PPP Subnets for Client Addresses" on page 149.
- If you're connecting two networks using SLIP or PPP, each network should have its own network number, and the systems that form the endpoints of the link should both run *routed*. You will find more information in "About Connected SLIP or PPP Networks" on page 149.

In all cases, the server should run the routing daemon, *routed*, unless you have a very small network and use static routing. For more information, see the *routed(1M)* and *route(1M)* reference pages. You can turn *routed* on using the *chkconfig* command:

```
# chkconfig routed on
```

Standalone clients should not run *routed*.

About Proxy-ARP Routing for SLIP Connections

If you have a small number of standalone hosts connecting to a SLIP server, you can use Proxy-ARP routing (proxy-ARP routing is not required for PPP connections, because PPP automatically installs ARP table entries if they are needed). In this system, each of the standalone hosts is assigned an Internet address from the server's network. Each client sets up a default route through the server, and the server advertises each of the client's addresses using the Address Resolution Protocol (ARP).

Setting Up Proxy-ARP Routing for SLIP Connections

To set up a default route through the server, you can use the *route* command after a SLIP connection is established. Add this command to the script that you use to start SLIP:

```
route add net default server-address 1
```

For each SLIP client, the server should issue an *arp* command:

```
arp -s client-hostname server-ethernet-address pub
```

Note that the server's Ethernet address is not the same as its IP address. To determine the server's Ethernet address, run the *arp* command from another system on the same Ethernet:

```
% arp dial-in.salad.com
dial-in.salad.com (192.70.79.7) at 8:0:69:9:4f:ef
```

Proxy-ARP Routing Example

The string of hexadecimal numbers separated by colons (8:0:69:9:4f:ef) is the server's Ethernet address.

The server can be set up to run the *arp* commands at boot time by placing the commands in a local network script. The following example shows a local network script that sets up ARP entries for a set of clients:

```
#!/bin/sh
#
# starting up local networking stuff
#

IS_ON=/etc/chkconfig
CONF=/etc/config
SERV_ADDR=8:0:69:9:4f:ef

if $IS_ON verbose ; then
    ECHO=echo
    VERBOSE=-v
else
    # For a quiet startup and shutdown
    ECHO=:
    VERBOSE=
fi

case "$1" in
'start')
    # setup proxy ARP for the dialin hosts
    # if this host has more than one interface,
    # you will need to hard-code the Ethernet MAC address
    # instead of letting it be determined at run time.
    # note that 4DDN (among others) may change the MAC address from
default!
    # and some AppleTalk packages change the output of `netstat
-ian`!
    arp -s client1 $SERV_ADDR pub
    arp -s client2 $SERV_ADDR pub
    arp -s client3 $SERV_ADDR pub
    ;;
'stop')
    # be nice and delete the ARP entries
    arp -d client1
    arp -d client2
    arp -d client3
```



```

        ;;
    *)
        echo "usage: $0 {start|stop}"
        ;;
esac
exit 0
#

```

Assume the preceding file was named */etc/init.d/network.local*, and you want it to run just after networking started, then you would use the following commands to create the startup and shutdown links:

```

ln -s /etc/init.d/network.local /etc/rc2.d/S31netlocal
ln -s /etc/init.d/network.local /etc/rc0.d/K39netlocal

```

Setting Up SLIP/PPP Subnets for Client Addresses

If you have a lot of standalone clients, it would be cumbersome to issue *arp* commands for all of them. A better strategy is to allocate client addresses from a special subnet set aside for SLIP and PPP clients. The server acts as a gateway to this net. The server must run *routed*, and it should use the **-F** option to reduce unnecessary network traffic. For example, if the SLIP/PPP subnet was 128.70.80, the string “-F 128.70.80” should be added to the */etc/config/routed.options* file.

The clients should set up default routes through the server, just as with proxy-ARP routing. In addition, if clients are to communicate with one another, they must have their Ethernet interfaces turned off:

```
% chkconfig network off
```

If clients have their Ethernet interfaces enabled, they will try to reach other clients on the same “net” through Ethernet, rather than through the server.

About Connected SLIP or PPP Networks

Connecting two networks using SLIP or PPP is probably the simplest case from a routing standpoint. Each network should have its own network number, and the addresses of the server and client should be allocated from their respective networks. Both client and server should run *routed*.

Routers on the main network may need to be configured to recognize and route to the new network.

Using Dynamic Address Allocation With PPP

If you are connecting to a service provider who uses PPP with dynamic address allocation, use the keywords **localhost=0,0** and **add_route** in the `/etc/ppp.conf` file, to allow the remote system to assign an IP address, and set up a default route through the remote system.

Configuring a Bidirectional Link

A simple SLIP or PPP link, in which one station must always initiate the connection, can be changed to allow either station to initiate the connection.

The first step is to establish a link in one direction. Once that link is working correctly, set up the link in the reverse direction. You should be able to use the same `ppp.conf` entry for both dial-in and dial-out.

Starting SLIP or PPP at Boot Time

To have a SLIP or PPP connection between networks start automatically, create a local network script, `/etc/init.d/network.local`, for that purpose. The script resides on the station where you want to initiate the link and should be linked to appropriate files in the `/etc/rc2.d` and `/etc/rc0.d` directories. For example, to automatically start the demand-dialed PPP link described in “Sample PPP Configuration for Dial-Out” on page 142, you would create a local network script on `caesar.salad.com`. The script should start PPP when called with the argument **start**, and terminate it when called with the argument **stop**.

```
#!/bin/sh
# ppp boot startup script
case $1 in
  start)
    /etc/killall ppp
    if /etc/chkconfig ppp && test -x /usr/etc/ppp -a -s /etc/ppp.conf
    then
```

```

        /usr/etc/ppp -r salad &
    fi
    ;;
stop)
    /etc/killall -TERM ppp
    ;;
*)
    echo "usage: $0 {start|stop}"
    ;;
esac

```

This script should be linked to the appropriate filenames in the `/etc/rc2.d` and `/etc/rc0.d` directories:

```

ln -s /etc/init.d/network.local /etc/rc2.d/S31netlocal
ln -s /etc/init.d/network.local /etc/rc0.d/K39netlocal

```

About Demand Dialing

If you have a SLIP or PPP link that is used irregularly but frequently, it is likely to be economical for you to make that link a demand-dialing link. With demand dialing, the system makes the telephone connection as needed, when there is network traffic to be transmitted, and it drops the connection when there is no traffic. By default, much of the non-essential network traffic does not cause a link to be established. For example, the traffic generated by the time daemon (*timed*) would not cause a call to be placed, but a request for file transfer would cause the connection to be made.

Setting Up Demand Dialing

To use demand-dialing, add the `-q` option to your *slip* command on the system that initiates the connection. Demand-dialing mode is also known as “quiet” mode. For complete information on this and other options, see the *slip(1M)* reference page.

PPP can be configured to use demand dialing by adding the “quiet” keyword to the *ppp.conf* file.

When using demand-dialing, it makes sense to start SLIP or PPP automatically at boot time, as described in “Starting SLIP or PPP at Boot Time” on page 150.

Note that the link is initiated only when there is traffic on the initiating end of the link.

NFS Over SLIP or PPP

You can run NFS over a SLIP or PPP link. NFS will be very slow because of the amount of information transferred in NFS transactions. You may be able to improve performance with these measures:

- Use NFS with modems faster than 9600 bps.
- Use one of the SLIP header prediction and compression options.

For more information regarding the SLIP compression options, **comp** and **cslip**, see the `slip(1M)` reference page.

- Adjust the NFS options **rsize**, **wsize**, **timeo**, and **retrans** when mounting NFS file systems.

To improve performance, read and write smaller blocks and specify longer timeouts. See the `fstab(4)` reference page for more information on NFS filesystem options.

File Transfer Over SLIP or PPP

File transfer over a serial link may be slow if other demanding utilities share the link. For faster file transfer, try using *uucp*. Using *uucp* is effective if you do not want to share the line with other utilities.

Troubleshooting SLIP and PPP Links

If your SLIP or PPP link seems to be connecting, but you can't reach systems on the remote network, you might have a routing problem. Try to reach the remote system with *ping*:

```
% ping -c 10 dial-in.salad.com
```

If the connection is working, you should see output something like this:

```

PING dial-in.salad.com (128.70.79.52): 56 data bytes
64 bytes from 128.70.79.52: icmp_seq=0 ttl=255 time=2 ms
64 bytes from 128.70.79.52: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=2 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=3 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=4 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=5 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=6 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=7 ttl=255 time=1 ms
64 bytes from 128.70.79.52: icmp_seq=8 ttl=255 time=2 ms
64 bytes from 128.70.79.52: icmp_seq=9 ttl=255 time=1 ms

```

```

----dial-in.salad.com PING Statistics----
10 packets transmitted, 10 packets received, 0% packet loss

```

If the connection is not working, you should see output like this:

```

PING dial-in.salad.com (128.70.79.52): 56 data bytes

```

```

----dial-in.salad.com PING Statistics----
10 packets transmitted, 0 packets received, 100% packet loss

```

If you can contact the remote host, but not other systems on the network, you probably have a routing problem. Check that your routing is set up as described in “SLIP and PPP Routing and Address Allocation” on page 146.

If you aren’t getting any connection at all, test the line with another utility. If you are familiar with *uucp*, you may want to establish a *uucp* link between the stations as a means of testing the connection. Most users will find that *cu* is an easier way to debug the link.

Note: *cu* requires a *direct* entry in the */etc/uucp/Devices* file. Refer to “UUCP Devices File” on page 211 for more details.

When debugging a SLIP connection, check each station separately. First check the port and modem on the local station by using a *cu* command like this:

```

cu -d -s speed -l port

```

For example, to test the port and modem installed on the station tuna.salad.com, you would use this *cu* command:

```
cu -d -s 38400 -l ttyf2
```

(It may be necessary to turn off the *(uu)etty* first by changing **respawn** to **off** on the line for the port in the file */etc/inittab*.)

The modem should respond. Many modems respond by printing *AT*. If the modem does not respond as expected, review the SLIP configuration procedure for the local station and review the modem configuration and documentation. If the modem does respond as expected, disconnect from *cu* by typing a tilde followed by a dot:

```
~.
```

Connect the stations you want to link as you would for the SLIP link. On the local station, use *cu* to call the remote station through the port and connection you have already verified with *cu*.

Check the connection to the remote station with a *cu* command like this:

```
cu -d -sspeed telno
```

For example, to test the connection between tuna and dial-in.salad.com, you would call dial-in from tuna with this *cu* command:

```
cu -d -s38400 5552002
```

You should see the local station tell the modem to call the remote station. Eventually, you should see the login prompt. Type the *send* strings from the */etc/uucp/Systems* file in response to the *expect* strings.

BIND Name Server

The Berkeley Internet Name Domain (BIND) server implements the Internet Domain Name Service (DNS) for the IRIX operating system. A name server is a network service that enables clients to name resources or objects in the network and share this information with other network objects. In effect, a name server is a distributed database system for objects in a computer network. All IRIX network programs can use BIND to store and retrieve station names and addresses. You can use BIND to replace the original *host* table lookup of information in the */etc/hosts* file.

BIND has two parts: the name server program, *named*, and a set of C library “resolver” routines that access the server. To set up as a name server you will have to install *named*. See “BIND Configuration Files” on page 161 for details. *named* is a daemon that runs in the background and responds to UDP and TCP queries on a well-known network port. The library routines reside in the standard C library, *libc.a*. The host-address lookup routine *getaddrinfo* and *sethostent* use the resolver routines to query the name server. The *getaddrinfo* function is defined for protocol-independent nodename-to-address translation. It performs the functionality of *gethostbyname(3N)* and *getservbyname(3N)*, but in a more sophisticated manner. For more information, see the *getaddrinfo(3)* man page.

The resolver library routines described in *resolver* include routines that build query packets and exchange them with the name server.

Note: The first six sections of this chapter describe earlier versions of BIND. The final section describes BIND version 9. If you are running IRIX 6.5.22 or later, see the BIND version 9 information in “Introduction to BIND Version 9” on page 181 and *BIND 9 Administrator Reference Manual* in */usr/share/doc/bind*.

The following topics are covered in detail in this chapter:

- “About Domain Name Service” on page 156
- “BIND Servers and Clients” on page 158

- “BIND Configuration Files” on page 161
- “Setting Up a BIND Configuration” on page 168
- “Managing the BIND Environment” on page 177
- “Debugging named” on page 179
- “Introduction to BIND Version 9” on page 181

About Domain Name Service

Host-table lookup routines, such as those using the */etc/hosts* file, require that the master file for the entire network be maintained at a central location by a few people. This approach works well for small networks where there are only a few stations and there is cooperation among the different organizations responsible for them. However, this approach does not work well for large networks where stations cross organizational boundaries.

The Domain Name Service eliminates the need for a single, centralized clearinghouse for all names. The authority for this information can be delegated to the organizations on the network that are responsible for it.

The Domain Name Service is organized as a hierarchical name space, like the IRIX file system. Figure 6-1 shows a small section of this hierarchy. Each subtree in the hierarchy is called a *domain*, and is given a label. At the top of the hierarchy is the root domain, labelled with the null label (“”). The name of the domain is the concatenation of all the domain labels from the root to the current domain. The labels are listed from right to left and are separated by dots. (For example, the name for the domain labelled fruit in Figure 6-1 would be fruit.salad.com.) A label must be unique only within its domain.

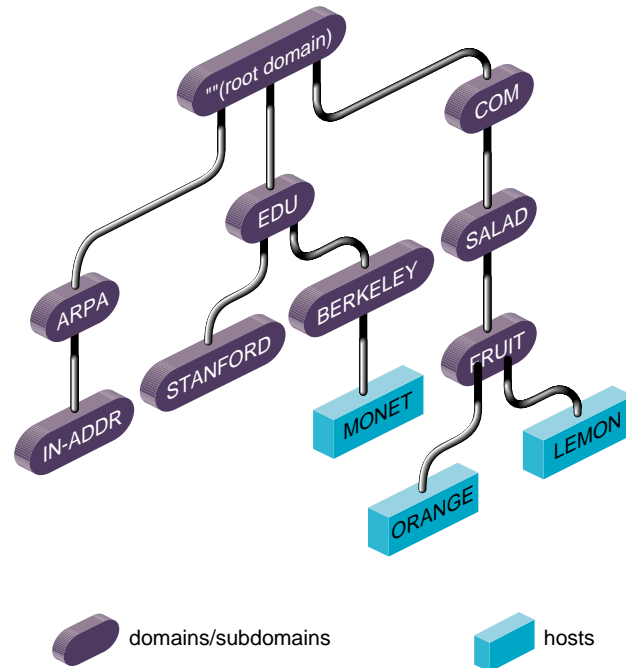


Figure 6-1 Partial View of Domain Name Space

Immediately below the root domain is a set of top-level domains. These top-level domains are relatively static, and are administered by the Network Information Center. These are the current top-level domains registered with the Network Information Center:

arpa	A temporary domain for stations, also used as the top-level domain for address-to-name mapping
com	Companies and businesses
edu	Universities and other educational institutions
gov	Government agencies
mil	Military organizations
net	Various network-type organizations and network management-related organizations, such as information centers and operations centers
org	Technical support groups, professional societies, or similar organizations

There are also many national domains, such as DE for Germany and FR for France.

The whole space is partitioned into several non-overlapping areas of authority called *zones*. Information in each zone is handled by the zone's "authoritative" or "master" name server(s). Each zone starts at a domain and extends down to the leaf domains, or to domains where other zones start. Zones usually represent administrative boundaries.

An example of a domain name for a station at the University of California, Berkeley is monet.berkeley.edu.

The top-level domain for educational organizations is edu. Berkeley is a subdomain of edu, and monet is the name of the station.

BIND Servers and Clients

BIND is based on a server-client relationship. There are several different classes of servers, with varying degrees of authority. This section discusses the interaction between various types of servers, and between servers and clients. This information is summarized in Table 6-1.

Table 6-1 summarizes the general characteristics for various BIND server configurations.

Table 6-1 BIND Server Configurations

Primary Server	Secondary Server	Caching-only Server	Forwarder Server	Slave Server
Authoritative server for the domain	"Delegated" authority from primary server	Non- authoritative	Non- authoritative	Non- authoritative
Loads data from local file	Loads data from primary server	Answers queries or forwards queries to authoritative servers	Answers recursive requests or interacts with other name servers before answering requests	Accesses data from specified list of servers (forwarders)

All server configurations must run the *named* server daemon. The *named* daemon is started automatically during station startup if the configuration flag **named** is "on." See the chkconfig(1M) reference page for more details.

The client accesses data from the name servers specified in its *resolv.conf* file. It does not run the domain server, *named*.

Note: Starting with the IRIX 6.5 operating system, the Unified Name Service (UNS) runs the daemon *nsd*, and controls the resolve order of the BIND client with the */etc/nsswitch.conf* file. See “How UNS Works With BIND” on page 192 for a full discussion.

BIND Master Servers

A master server for a domain is the authority for that domain. This server maintains all the data corresponding to its domain. Each domain should have at least two master servers: a primary master, and a secondary master to provide backup service if the primary is unavailable or overloaded. A server can be a master for multiple domains, serving as primary for some domains and secondary for others.

A primary master server is a server that loads its data from a file on disk. This server can also delegate authority to other servers in its domain. A secondary master server is a server that is delegated authority and receives its data for a domain from a primary master server. At boot time, the secondary server requests all the data for the given domain from the primary master server. This server then periodically checks with the primary server to see if it needs to update its data.

Root servers are the master servers for the root and top-level Internet domains. They are listed in the *root.cache* file described in “BIND root.cache File” on page 165.

BIND Slave and Forwarding Servers

A slave server always forwards queries it cannot satisfy locally to a fixed list of forwarding servers, instead of interacting with the master name server for the root and other domains. There may be one or more forwarding servers, and they are tried in turn until the list is exhausted.

A slave-and-forwarder configuration is useful when you do not want all the servers at a given site to interact with the rest of the Internet servers. The stations might be administratively prohibited from having Internet access. To give the stations the appearance of access to the Internet domain system, the stations could be slave servers

to the forwarding server on the gateway station. The gateway server would forward the queries and interact with other name servers on the Internet to resolve each query before returning the answer. A benefit of using the forwarding feature is that the central station develops a more complete cache of information, which all the stations can take advantage of. The use of slave mode and forwarding is discussed further in “Setting Up a BIND Configuration” on page 168.

There are two main reasons to use forwarders. First, if your station does not have full network access, it cannot send IP packets to the rest of the network. Therefore, it must rely on a forwarder with access to the network. Second, the forwarder can see all queries as they pass through the server and, therefore, builds up a more complete cache of data than the cache in a typical station name server. In effect, the forwarder becomes a meta-cache from which stations can benefit, thereby reducing the total number of queries from that site to the rest of the network.

BIND Caching-Only Server

A caching-only server is not authoritative for any domain. It services queries and asks other servers, who have the authority, for needed information. The results of queries are cached, to reduce traffic to the authoritative server. Query responses include a *time-to-live* field, which indicates how long they should be cached for.

BIND Clients

A BIND client accesses the name servers that run on other stations in the network. The *named* server does not run on the client station.

BIND Configuration Files

This section discusses the various BIND configuration files. Examples of these files are provided in “Setting Up a BIND Configuration” on page 168.

In IRIX, the *named* database files are stored in the */var/named* directory, and *named* is not installed by default. You must select *eo.e.sw.named* from the installation media. To confirm whether *eo.e.sw.named* has been installed, type

```
versions eo.e.sw.named
```

A *README* file contains a short summary of the setup procedure and a list of official names for BIND clients and servers. Typically, servers are also clients. BIND clients require the */etc/resolv.conf* file.

The */var/named/Examples* subdirectory contains sample *named* database files. The files in the *Examples* directory should be used and changed to reflect your setup. These files use the record format described in Appendix A, “BIND Standard Resource Record Format”. The database files needed to set up your BIND environment are these:

- *named.boot*
- *root.cache*
- *named.hosts*
- *named.rev*
- *localhost.rev*

Note: If your network has more than one domain, incorporate the domain name as part of the *named.hosts*, *named.rev*, and *localhosts.rev* filenames when you create your versions of these files.

The number and configuration of the database files depend on the server type.

Table 6-2 summarizes which database files are required for each type of server.

Table 6-2 named Database Files

Filename	Primary Server	Secondary Server	Caching-Only Server	Forwarder Server	Slave Server
named.boot	required	required	required	required	required
localhosts.rev	required	required	required	required	required
named.hosts	required	N/A	N/A	N/A	N/A
named.rev	required	N/A	N/A	N/A	N/A
root.cache	required	required	required	required	required

BIND Boot File

The boot file is first read when *named* starts up. It tells the server what type of server it is, which zones it has authority over, and where to get its initial data. The default name of this file is */etc/named.boot*. The template for this file is called */var/named/Examples/named.boot.master* (for primary server) and *named.boot.slave* (for secondary server).

Note: The *named.boot* file applies to versions of BIND prior to the IRIX 6.5.22 release. For the IRIX 6.5.22 release or later, BIND version 9 is supported and the *named.boot* file is automatically converted to the *named.conf* file (with the new format) using a new utility script called */usr/sbin/confignamed*. For more information on BIND version 9, see “Introduction to BIND Version 9” on page 181.

To use a different file, create or modify the */etc/config/named.options* file with this entry:

```
-b other-bootfile-name
```

The recognized boot file structures are described in the subsections that follow.

Specifying a Directory in the BIND Boot File

The directory line specifies the directory in which the name server should run, allowing the other filenames in the boot file to use relative pathnames.

```
directory /var/named
```

This entry is required. It makes sure *named* is in the proper directory when you try to include files by relative pathnames with \$INCLUDE. It also allows *named* to run in a location that is reasonable for dumping core, if necessary.

Specifying a Primary Master in the BIND Boot File

The line in the boot file that designates a primary server for a zone looks like this:

```
primary Berkeley.EDU named.hosts
```

The first field specifies that the server is a primary one for the zone stated in the second field. The third field is the name of the file from which the data is read.

Specifying a Secondary Master in the BIND Boot File

The line for a secondary server is similar to that for the primary, except that it lists addresses of other servers (usually primary servers) from which the zone data is obtained. For example:

```
secondary Berkeley.EDU 128.32.0.10 128.32.0.4 ucbhosts.bak
```

The first field specifies that the server is a secondary master server for the zone stated in the second field. The two network addresses specify the name servers that are primary for the zone. The secondary server gets its data across the network from the listed servers. It tries each server in the order listed until it successfully receives the data from a listed server.

If a file name is present after the list of primary servers, data for the zone is saved in that file. When the server first starts, it loads the data from the backup file if possible, and consults a primary server to check that the zone information is still up to date.

Specifying a Caching-Only Server in the BIND Boot File

All servers should have a line like this one in the boot file to prime the name server's cache:

```
cache . root.cache
```

All listed cache files are read when *named* starts up. Valid values are reinstated in the cache, and the root name server information in the cache files is always used to handle initial queries.

The name server needs to know the servers that are the authoritative name servers for the root domain of the network. The *root.cache* file primes the server's cache with the addresses of these higher authorities. This file uses the Standard Resource Record format (or Master File format) described in detail in Appendix F.

You do not need a special line to designate that a server is a caching server. What denotes a caching-only server is the absence of authority lines, such as *secondary* or *primary*, in the boot file.

Specifying Forwarders in the BIND Boot File

Any server can make use of forwarders. For example, a server capable of processing recursive queries may try resolving queries on behalf of other stations. The *forwarders* command specifies forwarders by Internet address as follows:

```
forwarders    128.32.0.10    128.32.0.4
```

Specifying Slave Mode in the BIND Boot File

You can use slave mode if, because of limited network access, use of forwarders is the only way to resolve queries. You can also use slave mode if you wish to prevent the name server from using forwarders other than those listed. Slave mode is activated by the following command in the boot file:

```
slave
```

If you use *slave*, you must specify forwarders. In slave mode, the server forwards each query to each of the forwarders until an answer is found or the list of forwarders is exhausted.

BIND named.hosts File

This file contains the host-address database for your domain. It is required for primary servers.

BIND named.rev File

This file specifies the IN-ADDR.ARPA domain, which is used to translate IP addresses into names. Because Internet addresses do not fall within domain boundaries, this special domain was formed to allow inverse mapping. The IN-ADDR.ARPA domain for a station has four labels preceding it. These labels correspond to the four octets of an Internet address in reverse order. All four octets must be specified, even if an octet is zero.

For example, the Internet address 128.32.130.12 is located in the domain 12.130.32.128.IN-ADDR.ARPA. This reversal of the address allows for the natural grouping of stations in a network.

An IN-ADDR.ARPA domain can also represent a network. For example, if the ARPANET is network 10, there is a domain called 10.IN-ADDR.ARPA.

BIND localhost.rev File

This file specifies the IN-ADDR.ARPA domain of the local loopback interface's network address, 127.0.0.1. The address is better known as the *localhost* address. Many important network programs depend on the information in this domain. This file is required on all servers.

BIND root.cache File

This file, by default, contains the initial cache data for root domain servers. It is required, in one form or another, on all servers.

BIND /etc/config/named.options File

This file is optional. It is used during station startup and by the *named.restart* script. Specify command-line arguments for *named* in this file. See the *named(1M)* reference page for details on the options.

hoResolution With `/etc/resolv.conf`

The only configuration file required by BIND clients is the `/etc/resolv.conf` file. This file is read the first time `getaddrinfo` is called. The `resolv.conf` file has several functions:

- It defines the default domain or the default domain search list.
- It specifies the ordering of host resolution services used by `getaddrinfo`.
- It lists Internet addresses of name servers.

The first two items apply to both client and server stations. The last item is required only by client stations. The file's format is described in detail in the `resolver(4)` reference page.

To set up a station as a client of remote servers, add **nameserver** entries for the Internet addresses of the name servers to `/etc/resolv.conf`. For example:

```
nameserver 128.32.130.12
```

You can specify up to three `nameserver` entries. It is usually not necessary to create this file if you have a local server running. An entry for the local server should use an Internet address of 0 (meaning "this station").

On client and server stations, the name in `/etc/sys_id` should be set to the fully qualified domain name. For example:

```
monet.Berkeley.EDU
```

However, if you choose not to use fully qualified domain names, add a line with the keyword `domain` and the station's domain to the `resolv.conf` file. For example:

```
domain berkeley.edu
```

The `getaddrinfo` library routine normally access station information in this order:

1. NIS
2. BIND
3. Local `/etc/hosts` file

Beginning with IRIX 6.5, this order is determined by the `hosts` keyword in `/etc/nsswitch.conf`. See "How UNS Works With BIND" on page 192 for details.

To enable the system manager to copy files from another station when it is in single-user mode, the */etc/hosts* file should contain entries for important stations in addition to the entries for the local station's network interface(s) and *localhost*. See the *hosts.equiv(4)* reference page for more information about the format.

Setting Up a BIND Configuration

This section provides an example of how a BIND environment might be organized and describes the procedure for configuring the various servers and client stations. The example assumes you are connected to the Internet. When setting up your own environment, replace the variables in the example with your own BIND environment variables. The example is based on these variables:

- The domain is *fruit.com*, network address 128.70.10, and the network is attached to the Internet.
- Primary server is *apples.fruit.com*, internet address 128.70.10.1, and specific configuration steps are detailed in “Configuring the Primary BIND Server” on page 169.
- Secondary server is *oranges.fruit.com*, internet address 128.70.10.2, and specific configuration steps are detailed in “Configuring the Secondary BIND Server” on page 173.
- Forwarding server is *banana.fruit.com*, internet address 128.70.10.3, and specific configuration steps are detailed in “Configuring the Forwarding BIND Server” on page 175.
- Caching-only server is *guava.fruit.com*, internet address 128.70.10.4, and specific configuration steps are detailed in “Configuring a Caching-Only BIND Server” on page 174.
- Slave servers are *pineapple1.fruit.com* and *pineapple2.fruit.com*, Internet addresses 128.70.10.8 and 128.70.10.9, and specific configuration steps are detailed in “Configuring a Slave BIND Server” on page 176.
- Clients are *plum1.fruit.com*, *plum2.fruit.com*, and *plum3.fruit.com*, Internet addresses 128.70.10.5, 128.70.10.6, and 128.70.10.7, and specific configuration steps are detailed in “Configuring the BIND Client” on page 177.

Note: Current versions of BIND do not support _ (underscore) in any component of the system or domain name.

Figure 6-2 illustrates the example BIND environment described above. Station names in the figure are shortened for illustrative purposes only.




Figure 6-2 Example BIND Configuration

Configuring the Primary BIND Server

Use this procedure to configure a primary server:

1. Log in as `root`.
2. Move to the named example directory:


```
cd /var/named/Examples
```
3. Copy the template files to the `/var/named` directory:


```
cp named.boot.master root.cache named.hosts \
named.rev localhost.rev /var/named
```

4. Move *named.boot.master* to the default filename:

```
cd ..
mv named.boot.master named.boot
```

5. Modify *named.boot* with your editor of choice to resemble the following:

```
;
; Boot file for apple.fruit.com, primary for fruit.com
;
directory /var/named
;type    domain    source host/file    backup file
cache    .          root.cache
primary  fruit.com  fruit.named.hosts
```

6. Modify the *named.hosts* file, here called *fruit.named.hosts*, to resemble the following:

```
; Authoritative data for fruit.com
;
@    IN    SOA    apples.fruit.com. named-mgr.apples.fruit.com.
      (1994021501 ; Serial
      10800      ; Refresh 3 hours
      3600       ; Retry 1 hour
      3600000    ; Expire 1000 hours
      86400 )    ; Minimum 24 hours
; authoritative name servers for fruit.com
      IN    NS    apples.fruit.com.
      IN    NS    oranges.fruit.com.
; address records for all hosts on the net
      IN    A     128.70.10.1
apples  IN    A     128.70.10.1
oranges IN    A     128.70.10.2
banana  IN    A     128.70.10.3
guava   IN    A     128.70.10.4
plum1   IN    A     128.70.10.5
plum2   IN    A     128.70.10.6
plum3   IN    A     128.70.10.7
pineapple1 IN  A     128.70.10.8
pineapple2 IN  A     128.70.10.9
localhost IN  A     127.0.0.1
; canonical or alias name for localhost
loghost IN          CNAME localhost
```

7. Modify the *localhost.rev* file to resemble the following:

```

;localhost.rev  -- PTR record for 127.1
;
@ IN SOA apples.fruit.com. named-mgr.apples.fruit.com.
    (1994021501 ;Serial
    10800      ;Refresh 3 hours
    3600       ;Retry  1 hour
    3600000    ;Expire 1000 hours
    86400 )    ;Minimum 24 hours
; authoritative name servers for fruit.com
    IN      NS      apples.fruit.com.
    IN      NS      oranges.fruit.com.
0 IN      PTR      loopback.fruit.com.
1 IN      PTR      localhost.

```

8. Modify the *named.rev* file (*fruitnamed.rev*) to resemble the following:

```

;
;      @(#)named.rev  1.1      (Berkeley)      86/02/05
;
@ IN SOA apples.fruit.com. named-mgr.apples.fruit.com.
    (1994021501 ; Serial
    10800 ; Refresh 3 hours
    3600 ; Retry  1 hour
    3600000 ; Expire  1000 hours
    86400 ); Minimum 24 hours
;authoritative name servers for fruit.com
    IN      NS      apples.fruit.com.
    IN      NS      oranges.fruit.com.
;named.rev addresses, by default, are the last two numbers
;of the internet addresses in reverse order, if Class B
;address.  If Class C address, then it's the last number.
1          IN      PTR      apples.fruit.com.
2          IN      PTR      oranges.fruit.com.
3          IN      PTR      banana.fruit.com.
4          IN      PTR      guava.fruit.com.
5          IN      PTR      plum1.fruit.com.
6          IN      PTR      plum2.fruit.com.
7          IN      PTR      plum3.fruit.com.
8          IN      PTR      pineapple1.fruit.com.
9          IN      PTR      pineapple2.fruit.com.

```

9. Use the default *root.cache* file if the primary server is attached to the Internet. If practical, you should obtain the most up-to-date list from *rs.internic.net* using anonymous FTP. The list is kept in the file *domain/named.root*.

```

;      This file holds the information on root name servers needed to
;      initialize cache of Internet domain name servers
;      (e.g. reference this file in the "cache . <file>"
;      configuration file of BIND domain name servers).
;
;      This file is made available by InterNIC registration services
;      under anonymous FTP as
;          file          /domain/named.root
;          on server     FTP.RS.INTERNIC.NET
;      -OR- under Gopher at  RS.INTERNIC.NET
;          under menu     InterNIC Registration Services (NSI)
;          submenu       InterNIC Registration Archives
;          file           named.root
;
;      last update:      Sep 1, 1995
;      related version of root zone:  1995090100
;
;      formerly NS.INTERNIC.NET
;
.           3600000  IN  NS      A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.  3600000      A      198.41.0.4
;
;      formerly NS1.ISI.EDU
;
.           3600000      NS      B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET.  3600000      A      128.9.0.107
;
;      formerly C.PSI.NET
;
.           3600000      NS      C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET.  3600000      A      192.33.4.12
;
;      formerly TERP.UMD.EDU
;
.           3600000      NS      D.ROOT-SERVERS.NET.
D.ROOT-SERVERS.NET.  3600000      A      128.8.10.90
;
;      formerly NS.NASA.GOV
;
.           3600000      NS      E.ROOT-SERVERS.NET.

```



```

E.ROOT-SERVERS.NET.      3600000      A      192.203.230.10
;
; formerly NS.ISC.ORG
;
.                          3600000      NS      F.ROOT-SERVERS.NET.
F.ROOT-SERVERS.NET.      3600000      A      39.13.229.241
;
; formerly NS.NIC.DDN.MIL
;
.                          3600000      NS      G.ROOT-SERVERS.NET.
G.ROOT-SERVERS.NET.      3600000      A      192.112.36.4
;
; formerly AOS.ARL.ARMY.MIL
;
.                          3600000      NS      H.ROOT-SERVERS.NET.
H.ROOT-SERVERS.NET.      3600000      A      128.63.2.53
;
; formerly NIC.NORDU.NET
;
.                          3600000      NS      I.ROOT-SERVERS.NET.
I.ROOT-SERVERS.NET.      3600000      A      192.36.148.17
; End of File

```

10. Enable *named* and reboot the station with the following commands:

```

chkconfig named on
reboot

```

Configuring the Secondary BIND Server

Use this procedure to configure a secondary server:

1. Log in as *root*.
2. Move to the *named* example directory:

```

cd /var/named/Examples

```
3. Copy the template files to the */var/named* directory:

```

cp named.boot.slave root.cache localhost.rev /var/named

```
4. Move *named.boot.slave* to the default filename:

```

cd ..

```

```
mv named.boot.slave named.boot
```

5. Modify *named.boot* to look like the following:

```
more named.boot  
  
;  
; Boot file for orange.fruit.com, secondary for fruit.com  
;  
directory /var/named  
; type      domain      source host/file  backup file  
cache      .              root.cache  
secondary  fruit.com  128.70.10.1      fruithosts.bak
```

6. Use the same *localhost.rev* file you installed on your primary server.
7. Use the same *root.cache* file you installed on your primary server.
8. Enable *named* and reboot the station with the following commands:

```
chkconfig named on  
  
reboot
```

Configuring a Caching-Only BIND Server

Use this procedure to set up a caching-only server:

1. Log in as *root*.
2. Move to the *named* example directory:

```
cd /var/named/Examples
```
3. Copy the template files to the */var/named* directory:

```
cp named.boot.master root.cache /var/named
```
4. Move *named.boot.master* to the default filename:

```
cd ..  
  
mv named.boot.master named.boot
```

5. Modify *named.boot* to look like the following:

```
more named.boot
;
;Boot file for guava.fruit.com,caching-only server for
;fruit.com
;Note that there should be one primary entry for each SOA
;record.
;
;
directory          /var/named
;type    domain    source host/file  backup file
cache    .          root.cache
```

6. Use the same *localhost.rev* file you installed on your primary server.
7. Use the same *root.cache* file you installed on your primary server.
8. Enable *named* and reboot the station with the following commands:

```
chkconfig named on
reboot
```

Configuring the Forwarding BIND Server

Use this procedure to set up a forwarding server:

1. Log in as *root*.
2. Move to the *named* example directory:

```
cd /var/named/Examples
```

3. Copy the template files to the */var/named* directory:

```
cp named.boot.master root.cache localhost.rev /var/named
```

4. Move *named.boot.master* to the default filename:

```
cd ..
mv named.boot.master named.boot
```

5. Modify *named.boot* to look like the following:

```
more named.boot

;
;Boot file for banana.fruit.com, forwarder server
;for fruit.com
;Note that there should be one primary entry for each
;SOA record.
;
;
directory          /var/named

;type      domain          source host/file    backup file
cache      .               root.cache
forwarders 128.70.10.1     128.70.10.2
```

6. Use the same *localhost.rev* file you installed on your primary server.
7. Use the same *root.cache* file you installed on your primary server.
8. Enable *named* and reboot the station with the following commands:

```
chkconfig named on

reboot
```

Configuring a Slave BIND Server

1. Log in as *root*.
2. Move to the *named* example directory:
3. Copy the template files to the */var/named* directory:

```
cd /var/named/Examples

cp named.boot.slave root.cache localhost.rev /var/named

4. Move named.boot.master to the default filename:

cd ..

mv named.boot.slave named.boot
```

5. Modify *named.boot* to look like the following:

```

;
;Boot file for pineapple1.fruit.com, slave server for
;fruit.com
;
directory          /var/named
;type              domain          source host/file    backup file
cache              .               root.cache
forwarders        128.70.10.3
slave

```

6. Use the same *localhost.rev* file you installed on your primary server.
7. Use the same *root.cache* file you installed on your primary server.
8. Enable *named* and reboot the station with the following commands:

```

chkconfig named on
reboot

```

Configuring the BIND Client

Use this procedure to set up a BIND client:

1. Log in as *root*.
2. Create or modify the *resolv.conf* file to include the default domain name, the host resolution order, and the list of name servers. It should look something like this:

```

domain fruit.com
nameserver 128.70.10.4
nameserver 128.70.10.2
nameserver 128.70.10.1
hostresorder bind local

```

3. Rebooting the client is suggested, but not required.

Managing the BIND Environment

This section describes the steps involved in maintaining the databases. It details how to add and delete a station from the domain and how to add a new subdomain. It also discusses some of the scripts that manage the BIND database.

Adding a New BIND Station

To add a new station to your zone files:

1. Edit the appropriate zone file for the station's domain.
2. Add an *A* record for each address of the station.
3. Add *CNAME*, *HINFO*, *WKS*, *RP*, and *MX* records (optional).
4. Add the reverse *IN-ADDR* entry for each station address in the appropriate zone files for each network the station is on.

Deleting a BIND Station

To delete a station from the zone files:

1. Remove all the station's resource records from the zone file of the station's domain.
2. Remove all the station's *PTR* records from the *IN-ADDR* zone files for each network the station was on.

Adding Another BIND Domain

To add a new subdomain to your domain:

1. Set up the other domain server, the new zone file, or both.
2. For each server of the new domain, add an *NS* record to the zone file of the parent domain.
3. Add any necessary glue address records. See Appendix A, "BIND Standard Resource Record Format" for details about glue records.

named Reload Script

This shell script sends the HUP signal to *named*, which causes it to read *named.boot* and reload the database. All previously cached data is lost. Use this script when *named* is running and you want the internal database for *named* to reflect any changes you have made. The `/usr/sbin/named.reload` script is an easy way to do this.

named Restart Script

This shell script terminates the running *named* and starts a new one. Use this script when you have made changes to the *named.boot* file, or whenever you need to place the server in a known state. The `/usr/sbin/named.restart` script is an easy way to restart *named*.

Debugging named

When *named* is running incorrectly, first check `/var/adm/SYSLOG` for any messages

Note: Most of these syslog messages are not valid for Bind Version 9. For more information, see “Introduction to BIND Version 9” on page 181.

For additional information, send *named* one of the following signals, using *killall(1M)* and defining *SIG* as *INT*, *ABRT*, *USR1*, or *USR2*:

```
/etc/killall -SIG named
```

INT	Dumps the current database and cache to <code>/var/tmp/named_dump.db</code> . This dumping should indicate whether the database was loaded correctly.
ABRT	Dumps statistics data into <code>/var/tmp/named.stats</code> . Statistics data is appended to the file.
USR1	Turns on debugging. Each subsequent <i>USR1</i> increases the debug level. There are 10 debug levels, and each prints more detailed information. A debug level of 5 is useful for debugging lookup requests. The output goes to <code>/var/tmp/named.run</code> .
USR2	Turns off debugging completely.

SYSLOG Error Messages

Using *syslog*, *named* logs certain errors to */var/adm/SYSLOG*. This section lists important error messages and their meanings.

- *cname* has CNAME and other illegal data
An alias has more than just a CNAME record. For example, since only “monet” should have the A record, the following is wrong:

```
ucbmonet IN CNAME monet
ucbmonet IN A 128.32.0.1
```
- Attempted to query myself on *ipaddr* as name server for *cname*
The station is listed incorrectly as a forwarder in the *named.boot* file.
- *zoneref: Masters for secondary zone cname unreachable*
This station is a secondary server for *cname*. The primary server returned an invalid data response or, because of a network problem, the station was not able to contact the primary server to obtain the current state of the zone information.
- *Lame delegation to cname1 received from ipaddr(purported server for cname2) on query on name [cname3]*
The message indicates that the remote server at the specified address is supposed to be authoritative for the *cname2* domain but has returned an indication that implies it is not. The remote server or the server for its parent domain is misconfigured. This error message can be disabled if *named* is started with the **-L lamedel** option.
- *MAXQUERIES exceeded, possible data loop in resolving cname*
The name server has tried to query too many other servers for the specified record. This might happen if each of two remote servers reply that the other has the desired information for *cname*.
- *Malformed response from ipaddr*
The remote DNS server at the specified address returned a malformed packet. This message typically indicates an error in the remote server.
- *Bogus root NS cname1 received from ipaddr on query on name[cname2] -- rejected*
The name server at the specified address improperly returned NS records for the root domain. The records have been ignored. You can disable this error message if *named* is started with **-L rootns**.

- Root NS *dname1* received from *ipaddr* on query on name [*dname2*]

The name server at the specified address returned NS records for the root domain. You can disable this informational message if *named* is started with the **-L rootns** option.

Debugging Name Servers With the nslookup Command

The *nslookup* command is a useful debugging tool for querying local and remote name servers. *nslookup* operates in either interactive or non-interactive mode. Interactive mode allows the user to query the name server for information about various stations and domains or print a list of stations in the domain. Non-interactive mode is used to print just the name and requested information for a station or domain. The following example of *nslookup* output gets the address record for the station monet.berkeley.edu:

```
Default Server:  ucbvax.berkeley.edu
Address:  128.32.133.1
> monet
Server:  ucbvax.berkeley.edu
Address:  128.32.133.1
Name:    monet.berkeley.edu
Address:  128.32.130.6
```

To exit, press **Ctrl+D** or enter **exit**. The *help* command summarizes available commands. Useful *nslookup* command options are **set type=Any**, **set type=Mx** and **set debug**. The complete set of commands is described on the nslookup(1C) reference page.

Introduction to BIND Version 9

With the IRIX 6.5.22 release, the name server provided by BIND has been upgraded to version 9.2.2. The majority of the features of BIND version 9.2.2 are implemented on IRIX except when noted below. This chapter provides an overview of BIND version 9.2.2.

Detailed information is available at the following location: <http://www.isc.org>. Follow the link to BIND documentation.

The BIND Standard Resource Record Format defined for BIND version 6 is not documented for BIND version 9. This information is still available in the Appendix A of this manual.

Note: Please note that AAAA Record Format is the same as A Record Format except that the address has an IPv6 format.

BIND Version 9 is described in the following sections:

- “BIND Version 9 Overview” on page 182
- “Migrating to BIND 9” on page 183
- “Implementation differences for BIND 9 on IRIX” on page 184

BIND Version 9 Overview

BIND version 9 is a major rewrite of nearly all aspects of the underlying BIND architecture. Some of the important features of BIND 9 are, as follows:

- DNS security
 - DNSSEC (signed zones) - not supported on IRIX
 - TSIG (signed DNS requests)
- IP version 6
 - Answers DNS queries on IPv6 sockets
 - IPv6 resource records (A6, DNAME, and so on.)
 - Bitstring labels
 - Experimental IPv6 Resolver Library - not supported on IRIX
- DNS protocol enhancements
 - IXFR, DDNS, Notify, and EDNS0
 - Improved standards conformance
- Views

One server process can provide multiple “views” of the DNS namespace, for example, an “inside” view to certain clients, and an “outside” view to others.
- Multiprocessor support
- Improved portability architecture

Migrating to BIND 9

This section describes issues you need to consider when migrating to BIND 9.

BIND 9 uses a new format for its configuration file called the `/etc/named.conf` file. When `named` is restarted (using the `/usr/sbin/namedc` start script upon system startup), the `/etc/named.conf` file is automatically created. The `named.boot` file is automatically converted to the `named.conf` file (with the new format) using a new utility script called `/usr/sbin/confignamed`. For more information, see the `confignamed -h` option for help on this command.

A new program called `rndc` (Name Server Control Utility) is used to communicate and perform operations with a running `named` process. For more information, see the `rndc(8)` man page. The `rndc` command is used by `namedc` startup script.

This utility requires a configuration file `/etc/rndc.conf` which is shared between the `named` and `rndc` commands. This file is also generated (if it does not exist) when `namedc start` is performed at system startup.

The `named.reload` and `named.restart` scripts are now wrappers to the `namedc` script and perform the same operations as before.

For the UDP queries, BIND 9 uses system assigned ports rather than port 53 which is used by BIND 4. This may conflict with some firewalls. The following directives in `/etc/named.conf` file allows you to specify a port to use:

```
query-source address * port 53;
transfer-source * port 53;
notify-source * port 53;
```

Zone files are the same format except for minor differences as described in the following paragraphs.

BIND 9 no longer uses the `minimum` field to specify the TTL of records without an explicit TTL. Use the `$TTL` directive to specify a default TTL before the first record without an explicit TTL, as follows:

```
$TTL 3600
@      IN      SOA      ns1.example.com. hostmaster.example.com. (
                                           2001021100
```

```
7200
1200
3600000
7200 )
```

BIND 9 does not support multiple CNAMEs with the same owner name. For example, the following is **not** supported:

```
www.example.com. CNAME host1.example.com
www.example.com. CNAME host2.example.com
```

BIND 9 does not support "CNAMEs with other data" with the same owner name, ignoring the DNSSEC records (SIG, NXT, KEY) that BIND 4 did not support. For example, the following is **not** supported:

```
www.example.com. CNAME host1.example.com
www.example.com. MX 10 host2.example.com
```

BIND 9 is less tolerant of errors in master files. Check your logs and fix any errors that are reported. The `named-checkzone` program can also be to check master files.

Outgoing zone transfers now use the `many-answers` format by default. This format is not understood by certain old versions of BIND 4. You can work around this problem using the `transfer-format one-answer;` option. Since these old versions all have known security problems, the correct fix is to upgrade the slave servers.

Implementation differences for BIND 9 on IRIX

In addition to the exceptions noted above, the current implementation of BIND 9 in IRIX does not include the BIND 9 resolver components. IRIX continues to use the Unified Name Service (see "How UNS Works With BIND" on page 192) and `nsd`. The client interaction with BIND 9 server-side code is compatible, except that this release does not support the A6 record, a recent part of IPv6, on IRIX clients.

Unified Name Service

The Unified Name Service (UNS) is a name service layer provided with the IRIX operating system to translate and simplify name service requests. Name servers provide network services that enable clients to name resources or objects in the network and share this information with other network objects.

In the past, each new name service was implemented in an application using code in the standard C library. As a new name service was added to a network, configuration files representing information about system resources and accounts were added along with a number of library routines. When the concept of distributed name space administration was conceived, the process became more complex.

To simplify this, the unification layer provided by UNS was developed. All IRIX networked programs can use UNS to implement other name services.

The following subjects are discussed in this chapter:

- “About Unified Name Service” on page 186
- “Overview of UNS Operations” on page 186
- “How UNS Works With NIS” on page 189
- “How UNS Works With BIND” on page 192
- “How UNS Works With NFS” on page 193
- “How UNS Works With LDAP” on page 194
- “Setting Up a UNS Configuration” on page 197
- “UNS Protocol Libraries” on page 197
- “Troubleshooting nsd” on page 203

About Unified Name Service

UNS translates the results of name service requests from a number of different protocols into a single file-based protocol. A protocol is a set of rules, data formats and conventions that determine how data is transferred between network components. A library is the implementation of a protocol. Among several standard protocols that the IRIX operating system provides are the protocols DNS and NIS which deal with name service requests. Given multiple protocols, a single overriding protocol simplifies name service requests.

Unified Name Service has three major components:

- the name service daemon *nsd*
- specific application programming interface (API) routines in the C library
- several protocol libraries

The name service API is left unchanged from previous releases of the IRIX operating system to maintain library-level compatibility. No applications should need to recompile to take advantage of UNS components.

Overview of UNS Operations

At system start up, the daemon *nsd* takes two actions: it initiates a filesystem namespace rooted at */ns*. It then reads the UNS configuration file */etc/nsswitch.conf* which specifies the resolve order of each of the supported tables, and protocols.




Figure 7-1 Action of the nsd Daemon With Name Service Protocols

The daemon *nsd* has no awareness of protocols itself, it simply sets up a front-end interface to the different name services and runs between them. If a request comes from NIS to resolve a password into a system name, *nsd* passes the request to the appropriate library based on the resolution of the request pathname. A request issued from NIS `/etc/passwd`, for example, is passed to `/ns/.local/.nis/passwd`, so that in effect UNS acts as an alias.

When the *nsd* daemon is called, it reads the configuration file `nsswitch.conf`, which contains a list of protocols and tables. A table is made up of searchable rows and columns. The hosts table consisting of `hosts.byname` and `hosts.byaddr`, effectively allows a search either by hostname or by host address.

The file system namespace started by *nsd* is rooted at */ns* and is used by the underlying interfaces to retrieve name system data. This is a dynamic file system, timed out after 30 seconds, so that if you issue the command *ls* from */ns* you may not see the files.

The results of any name search are held in the *.local* cache file in the */ns* namespace. The *.local* file is the local user's view of the namespace. When a file is accessed in the */ns* directory, an entry is added to a cache file under */var/ns/cache*, creating what is in effect a shadow file. There is a cache file for each of the tables supported by the name service daemon. For example, if a call comes to NIS, the NIS database files on the local host are searched first. This means network traffic is reduced and searches are swifter.




Figure 7-2 Partial View of Dynamic UNS Files

The format of these files mimics the format of the historic configuration files that they replace. Using *passwd* as an example, each file under the directories */ns/*/passwd.byname* is made up of lines separated by newline characters, of the format:

```
login:password:uid:gid:gecos:directory:shell
```


Since the file `/etc/nsswitch.conf` specifies the resolve order for each of the supported tables, including the host table, the `hostresorder` line in `/etc/resolv.conf` is ignored.

The following table shows the name services that have been supported as part of the operating system since IRIX 6.2 with their client and server services, together with the binding that links the client and the server:

Table 7-1 Protocols With Historically Supported Services

Protocol	Client	Binding (Linking)	Server
DNS	Resolver Library	resolv.conf	named
NIS		ypbind (now nsd)	ypserv (now nsd)
Files	getx by y() getx by ent()		

Protocols and tables provided with the initial UNS release are client-side DNS; client-side NIS; Files; MDBM, the NIS database files; NDBM; Berkeley DB; Nisserv, the replacement for `ypserv`; and LDAP, the lightweight alternative for X500. The operation of these libraries is explained in “UNS Protocol Libraries” on page 197.

The following sections explain various aspects of UNS:

- “How UNS Works With NIS” on page 189
- “About UNS and the NIS Database” on page 192
- “How UNS Works With BIND” on page 192
- “How UNS Works With NFS” on page 193
- “How UNS Works With LDAP” on page 194.

How UNS Works With NIS

The NIS network lookup service provides a centralized database of information about the network to systems participating in the service. Of the number of information sources provided to network applications, typically the default lookup order is NIS first, DNS (BIND) second, then the appropriate local files.

Unified Name Service (UNS) does not change the fundamental operation of NIS, although how the server daemon *ypserv*, and the binding daemon *ypbind*, work are changed.

Historically, the NIS daemon *ypserv* is the database server responsible for answering client inquiries and updating the database. *ypserv* runs only on NIS server machines with a complete NIS database. The NIS binder daemon *ypbind*, would run on all NIS clients and was responsible for remembering information necessary for communicating with *ypserv*. A system can function as a server and a client simultaneously, and can therefore be running both *ypbind* and *ypserv*.




Figure 7-3 Historical Operation of NIS Name Lookups

In earlier versions of the IRIX operating system, NIS daemons are started by the master network script `/etc/init.d/network`, when the NIS daemon flags are set “on.” Beginning with the IRIX 6.5 operating system, the UNS file `/etc/nsswitch.conf` controls the resolve order of the NIS client, and so overrides the previous actions of *ypserv*.




Figure 7-4 Operation of the nsd Daemon With NIS

The *nsd* daemon reads the UNS configuration file */etc/nsswitch.conf*, the default copy of which contains a line of this format:

```
hosts:nis dns files
```

In this instance, the NIS hosts table is referred to before the DNS or local files. The *nsd* daemon then implements the NIS protocol in a shared library stored in the directory */var/ns/lib*. This is explained further in “UNS Protocol Libraries” on page 197.

Further changes to NIS server utilities involve *ypinit*, *ypmake*, and *make.script*, which are now wrappers around a perl script called *mdbm_parse*. This script does the actual parsing, instead of a makefile and binaries. This does not change visible system behavior.

For details of the NIS environment, refer to *NIS Administration Guide*. For more about the activity of the *nsd* daemon, see “UNS Configuration File” on page 196.

About UNS and the NIS Database

The NIS database is composed of a group of files known as maps. Maps are composed of keys and values. For example, in the map called *hosts.byname*, the keys are the names of the individual systems.

Maps are created from input files, (usually standard ASCII files) and made into files in database record format. Historically these backend database files, called *dbm*, had output files with *.dir* and *.pag* extensions, each pair being a map. Files in DBM format are restricted both in the number of entries, and the relative size of each entry.

Database backend files in the UNS environment are called *mdbm*, and output files often have *.m* extensions. The database itself needs only one file so there are not separate *.dir* and *.pag* files. Files in MDBM format handle millions of entries; each entry can be up to 64 KB long (4 KB by default). This simpler single-key file format leads to faster response.

With UNS, the new location for the NIS backend data storage is */var/ns/domains*. Each table is still stored in multiple files, one per key. Keys that are looked up are cached into local hash files found in */var/ns/cache/*.

How UNS Works With BIND

The Berkeley Internet Name Domain (BIND) implements the Internet Domain Name Service (DNS) for the IRIX operating system. All IRIX network programs can use BIND to replace the original host table to look up information in the */etc/hosts* file.

BIND server configurations run the *named* server daemon, which is started automatically during station startup if the configuration flag *named* is set "on." See the *chkconfig(1M)* reference page for more details.

The resolve order on a typical BIND server is determined through the *resolv.conf* file, an example of which follows:

```
noodle% cat resolv.conf
search example1.com example2.com
#hostresorder nis bind local # now defunct under UNS
nameserver 192.99.89.54 #bindle
nameserver 127.0.0.1
```

The basic BIND client normally accesses data from the name servers specified in its own *resolv.conf* file. It does not run the domain server *named*.

Starting with the IRIX 6.5 operating system, the UNS file */etc/nsswitch.conf* controls the resolve order of the BIND client and overrides the resolve order of *resolv.conf*. These lines from the default */etc/nsswitch.conf* show the hosts table reference:

```
spanker% cat nsswitch.conf
```

```
.....
```

```
group: files nis
```

```
hosts: nis dns files
```

UNS provides DNS client capabilities.

In the UNS environment, the *nsd* daemon implements the DNS protocol in a shared library stored in the directory */var/ns/lib*. This is explained further in “UNS Protocol Libraries” on page 197. For further details on DNS, see Chapter 6, “BIND Name Server.”

How UNS Works With NFS

NFS is an application layer service that can be used on a network running the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP). With NFS a single copy of a file such as */usr/local* can be accessible by all systems on the network. Remote files appear to be “local” to each user’s system. The NFS server daemon *nfsd* runs on the server and accepts RPC calls from clients. NFS caches file attributes on the client side, known as the front-end filesystem, so that certain operations do not have to go all the way to the server.

When a file operation like *ls -l* is requested, the data inside the file is not touched. The information delivered is about the file’s attributes (for example, size, time of access, and owner), and those attributes are valid for some 3 seconds. Client changes are made on the locally cached copy, and if the file’s attributes remain unchanged for about 60 seconds, they are flushed from the cache and written back to the server. The server is known as the back-end filesystem and contains the definitive archived copy of the data.

Starting with the IRIX 6.5 operating system, the UNS environment provides the local files, which are stored in `/ns/.local`, the local user's view of the namespace. As with the other name services, when a file is accessed in the `/ns/.local` directory, a cache entry is added into a cache file under `/var/ns/cache`, creating what is in effect a shadow file. The `nsd` daemon presents the appearance of a mounted filesystem. There is actually no remote NFS server.

Details of the SGI implementation of the Sun Microsystems Open Network Computing Plus (ONC+) distributed services, previously referred `ldap.conf(4)` to as Network File System (NFS), are explained in *ONC3/NFS Administrators Guide*.

How UNS Works With LDAP

The Lightweight Directory Access Protocol (LDAP) is a simplified version of the X.500 Directory Access Protocol implemented on TCP/IP. One of LDAP's special attractions is that it allows you to define the database schema, specifying how information will be organized and retrieved from the database. Of course, an LDAP client must know of this organization in order to use the database, and this is done with the LDAP client configuration file, `/var/ns/ldap.conf`.

At system startup, the `nsd` daemon reads the configuration file `/etc/nsswitch.conf` which specifies the resolve order of each of the supported tables and protocols. If one of those listed is LDAP, the LDAP client configuration file `/var/ns/ldap.conf` is read.

In this file, each domain using the LDAP protocol library has a section specifying which LDAP servers to use, the parameters to search on each server, how to map each name service request to an LDAP search filter, and how to transform the response into a file format. With this flexibility you can use any schema for storing the configuration information. For technical information about LDAP, refer to RFC 1777. For details of the configuration file, see the `ldap.conf(4)` man page.

Namespace Format

The namespace is of the format */ns/domain/table/protocol/key*. The protocol directory is often left out. This results in all the protocols being searched in the order given in */etc/nsswitch.conf*. The *.local* domain represents the local system view of the namespace. A special file *.all* gives the full enumeration of the data for a table if supported by a protocol. Files and their purposes are enumerated in Table 7-2:

Table 7-2 UNS Files and Their Purposes

File Name	Purpose
<i>/ns/.local/passwd.byname/root</i>	Root entry in local passwd using whatever protocols are supported.
<i>/ns/.local/passwd.byname/.nis/root</i>	Root entry in local passwd using only the NIS protocol.
<i>/ns/.local/passwd.byname/.nis/.all</i>	All password entries in the NIS passwd table for the local domain.
<i>/ns/engr/passwd.byname/root</i>	Root entry is the engr domain passwd map using any available protocol.
<i>/ns/sgi.com/hosts.byname/sgi</i>	Finds the host address for the machine sgi in the sgi.com domain. (The local host must be a server for sgi.com.)
<i>/ns/.local/hosts.byname/sgi.sgi.com</i>	Finds the host address for the machine sgi.sgi.com in the hosts map. (Data may be anywhere.)

For example, to look up the password entry for the root user in the domain *engr.example1.com*, you would simply give this command:

```
# cat /ns/engr.example1.com/passwd.byname/root
```

The directory *.local* is created for the local domain so that the root password entry for the local domain can always be found in the file */ns/.local/passwd.byname/root*.

The file *.all* in each table directory enumerates the entire password table. If you need to list every password entry for the local domain using all the library routines listed in *nsswitch.conf*, give this command:

```
# cat /ns/.local/passwd.byname/.all
```

This gives you a large concatenation of every user in the local domain.

Finally, a special directory *.library* (for example *.nis*) is created under each table directory for each of the libraries listed for that table in *nsswitch.conf*.

The *ns_lookup()* library routine always opens files under the *.local* domain namespace mounted on */ns* to satisfy the requests from name service lookups, so this should not be changed.

UNS Configuration File

The sole UNS configuration file is */etc/nsswitch.conf*, and is made up of lines of the format:

```
map: library library library
```

For example:

```
hosts: nis dns files
```

This file specifies the libraries and the order in which they are to be used. If any path element does not exist, name service library routines are called by the *nsd* daemon in progressive order until the element is found.

The system administrator need only verify that *nsd* is running—see “Troubleshooting *nsd*” on page 203. Normally the *nsd* daemon is activated at system startup time from */etc/init.d/network* if the configuration flag is set “on” (see *chkconfig(1M)*). If the *nsd* configuration flag is not set “on”, the only name service supported is local files. The *nsd* daemon converts each name service request into a pathname.

When it is started, *nsd* creates the dynamic file system namespace rooted at */ns* which is used by the underlying interfaces to retrieve name system data.

The format of these files mimics the format of the historic configuration files that they replace. Using *passwd* as an example, each file under the directories */ns/*/passwd.byname* is made up of lines separated by new-line characters, of this format:

```
login:password:uid:gid:gecos:directory:shell
```

Since the file */etc/nsswitch.conf* specifies the resolve order for each of the supported tables, including the host table, the *hostresorder* line in */etc/resolv.conf* is ignored.

There is one *nsswitch.conf* file for each “domain” supported by this daemon. Every machine has a local domain, *.local* which has the configuration file in */etc/nsswitch.conf*. Server machines support multiple domains and have a configuration file in */var/ns/domains/domainname/nsswitch.conf* for each domainname.

Setting Up a UNS Configuration

The UNS configuration is automatic with the installation of the IRIX operating system. The *nsd* daemon runs and uses the */etc/nsswitch.conf* file to determine the configuration and protocols of the name services being requested. There is one *nsswitch.conf* file for each “domain” supported by this daemon. Every machine has a local domain *.local* which has the configuration file in */etc/nsswitch.conf*. Server machines support multiple domains and have a configuration file in */var/ns/domains/domainname/nsswitch.conf* for each domainname.

However, additional steps are needed to set up a system as a NIS Server:

1. Run *ypinit* to setup a system as a NIS server. It copies a default *nsswitch.conf* file into place and parses the system configuration files into the *mdbm* hash files. Refer to the *ypinit(1M)* reference pages for details.
2. Run *yppmake* to rebuild and distribute the NIS database. Refer to the *yppmake(1M)* reference page for details.
3. Run *mdbm_dump* to view the current contents of an *mdbm* hash database.

UNS Protocol Libraries

The protocol libraries provided with the initial UNS release are:

- Files—see “Files Library” on page 199.
- NIS, the Network Information Service—See “NIS Protocol Library” on page 199.
- DNS, the Domain Name Service—See “DNS Protocol Library” on page 199.
- MDBM, local hash files—See “MDBM Protocol Library” on page 199.
- NDBM, an earlier version of MDBM—See “NDBM Protocol Library” on page 200.
- Berkeley DB, Berkeley hash files —See “Berkeley DB Protocol Library” on page 200.

- Niserv, the replacement for ypserv —See “Niserv Protocol Library” on page 200.
- LDAP, the lightweight alternative for X500 —See “LDAP Protocol Library” on page 201.

As Figure 7-5 indicates, the *nsd* daemon directs requests to specific protocol libraries based on the resolution of a process pathname.

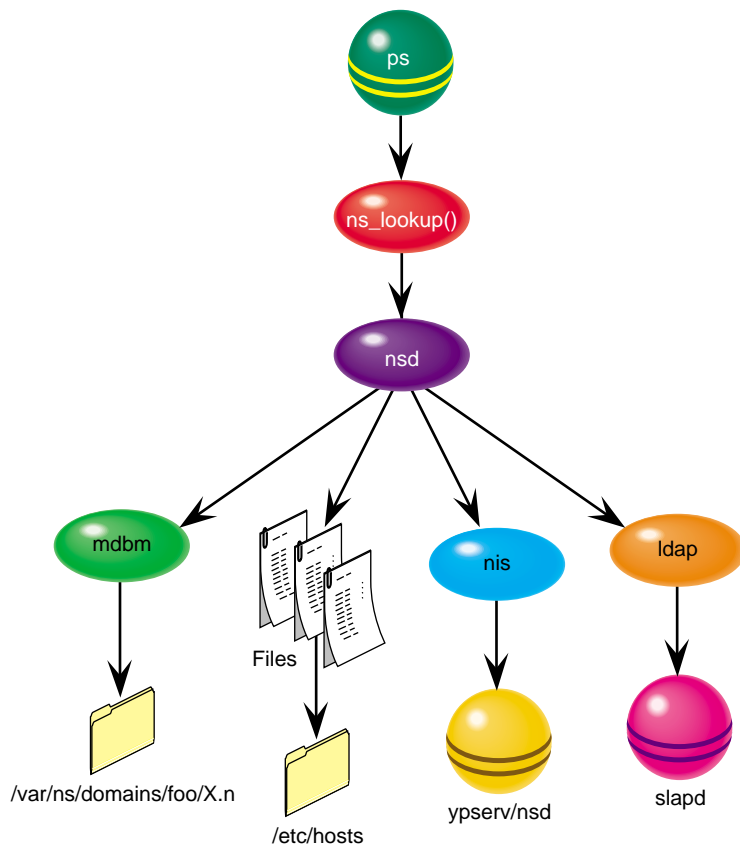


Figure 7-5 Action of *ns_lookup* in Selection of Protocol Library

Files Library

The files library parses historic format configuration files such as *passwd* and *hosts*. If looked up in the *.local* domain, the file in */etc* is opened and searched. If looked up in some other domain, then a file by the same name is opened under */var/ns/domains/domainname* for the requested domainname. For more about how these attributes function, see “UNS File Structure Attributes” on page 201. See also the reference page *files(7P)*.

NIS Protocol Library

The NIS protocol library makes a request to a remote *ypserv* daemon using version 2 of the NIS protocol. It also takes the place of the *ypbind* daemon on the IRIX operating system. If the NIS protocol is supported on any table, then the library creates a new service socket that answers NIS bind requests from the old style NIS API routines *yp_match()*, *yp_all()*, and so on.

If a key is looked up in the *.local* domain then the current domainname of the system as specified by the *domainname* command is used. This is typically set in the */var/yp/ypdomain* file. If the domain is unspecified then the server is dynamically located using the *ypbind* protocol. The maximum number of hops to use in multicast binding can be specified with the **nis_multicast** attribute which defaults to 32. See “Setting Domain Attributes” on page 201.

DNS Protocol Library

The DNS protocol library makes a request to a remote *named* daemon using the DNS protocol. The DNS library parses the file */etc/resolv.conf* to determine the default search path, nameserver list, and options. If a key is looked up in the *.local* domain then it will follow the search path; otherwise the specific domain is used. See “Setting Domain Attributes” on page 201 for a discussion on setting these attributes.

MDBM Protocol Library

The MDBM protocol library searches for data in *mdbm* hash database files on the local host. Names for these files are the same as the table with a *.m* extension added (for example: *hosts.byname.m*). For the *.local* domain these are found in the */etc*

directory, otherwise, they are found in the `/var/ns/domains/domainname` directory for each domainname. The flat files can be parsed into `mdbm` hash files using the `mdbm_parse` script provided with the `nisserv` library. See the `mdbm` reference page, and “Setting File Attributes” on page 202 for an explanation.

NDBM Protocol Library

The NDBM protocol library searches for data in `ndbm` hash database files on the local host. Names for these files are the same as the table with a `.pag` and `.dir` extension added (for example: `hosts.byname.pag` and `hosts.byname.dir`). For the `.local` domain these are found in the `/etc` directory, otherwise, they are found in the `/var/ns/domains/domainname` directory for each domainname. See the `ndbm(7P)` reference page, and “Setting File Attributes” on page 202 for an explanation.

Berkeley DB Protocol Library

The Berkeley DB protocol library searches for data in `db` hash database files on the local host. Names for these files are the same as the table with a `.db` extension added (for example: `hosts.byname.db`). For the `.local` domain these are found in the `/etc` directory, otherwise, they are found in the `/var/ns/domains/domainname` directory for each domainname. See the `berkeleydb(7P)` reference page, and “Setting File Attributes” on page 202 for an explanation.

Nisserv Protocol Library

The `nisserv` protocol library replaces the `ypserv` daemon in the IRIX operating system. This library will look up data in `MDBM` hash files on the local disk and provide them to remote machines using version 2 of the NIS protocol. The file naming rules are the same as those for `MDBM`. If the `nis_secure` attribute is set or the key `YP_SECURE` exists in the map file, then client systems are required to be using a port number less than 1024. This provides some weak security for maps with confidential information. A number of other special keys exist in the maps to provide the data for NIS protocol procedures. These keys include `YP_MASTER_NAME`, `YP_LAST_MODIFIED`, `YP_INPUT_FILE`, `YP_OUTPUT_NAME`, and `YP_DOMAIN_NAME`. These keys are historic from the previous implementation of NIS. For more information, see the `nisserv(7P)` reference page.

LDAP Protocol Library

The LDAP protocol library converts local name service requests into remote LDAP protocol packets, then formats the result so that it may be used by the standard name service API. All server information and formatting rules are read from a single configuration file `/var/ns/ldap.conf`. This file is read when the library is initialized or every time the name server daemon receives a SIGHUP signal.

UNS File Structure Attributes

The `nsd` daemon implements an in-memory filesystem. File structures represent the state for each file, and extended attributes are supported on each name service file. The attributes on the file depend on the library that looked them up, but always include: domain, table, key, timeout, source, version, and server. Attributes are inherited from parent directories. If an attribute such as **timeout** exists on a directory then all the files in that directory will use the same value for **timeout**, unless it is overridden for a particular file. The global attributes, timeout, hostname, and program can be set by command line arguments.

This section explains how various attributes are set or queried:

- “Setting Domain Attributes”
- “Setting Table Attributes”
- “Setting Library Attributes”
- “Querying Attributes”

Attributes to `nsd` are given either on the command line as `nsd -a key=value` or specified in the `nsswitch.conf` file. Following any change, you need to restart `nsd` by giving the command:

```
# killall -HUP nsd
```

Setting Domain Attributes

In the `nsswitch.conf` file, set the attributes on a domain by lines in this format:
(key1=value1, key2=value2)

Setting Table Attributes

Attributes on a table follow the table name in the file:

```
passwd(timeout=10): nis files
```

Setting Library Attributes

Attributes on a library (for the table) are given following the library name:

```
passwd: nis(timeout=10) files
```

Setting File Attributes

Attributes on a particular file are specified in the filename:

```
/ns/.local/passwd.byname/root(timeout=10)
```

Attributes can be changed by the library routines at will, and if an attribute is unspecified, then a default value is used.

Querying Attributes

Attributes can be listed or queried using the *attr* command:

```
attr -l /ns/.local/passwd.byname/root
```

```
attr -g timeout /ns/.local/passwd.byname/root
```

Cache Tuning

The size of the cache can be controlled with the parameters shown in Table 7-3, which are set in the *nsswitch.conf* file.

Table 7-3 Cache Tuning Parameters

Parameter	Description	Default/Source
size	Total maximum pages, squared	8
timeout	In seconds, length of time before cache expires	300
pagesize	In bytes to the power of 2	4 Kbytes
mode	Permissions, can be varied	0755
library	Name of the library that provided the data, as given in <i>nsswitch.conf</i>	Set by <i>xattr</i> command

Troubleshooting nsd

General Approach to Troubleshooting

Frequently the system log SYSLOG lists errors and points to solutions. The system log is available from the desktop. Error messages might take this form:

```
NFS3 server (local host) nsd not responding
NFS3 server local host (nsd) not responding
Stale NFS File Handle
```

Remote Service Errors

If you see consistent or random errors in remote services such as *ping* or *rlogin* errors, *nsd* may have failed. To check the status of *nsd*, enter:

```
ps -ef | grep nsd
```

If *nsd* does not appear in the list of processes running, restart *nsd*. Simply enter *nsd* at the prompt.

Mistakes in `nsswitch.conf`

A less common error message might be this:

```
no such protocol
```

In such a case, check for misspellings in `nsswitch.conf`.

Deciphering `nsd` Signals

Some signals `nsd` supports are familiar in other contexts. Send `nsd` the following command, defining **SIG** as **USR1**, **USR2**, **HUP**, or **TERM**.

```
kill -SIG nsd
```

- | | |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USR1 | Forces <code>nsd</code> to create a state file: <code>/usr/tmp/nsd.dump</code> which is a batch file reflecting the state of the file system at the moment. |
| USR2 | Increases the logging level. As you repeat the command, logging becomes more verbose. Beginning at 0, 6 is the most verbose. |
| HUP | Is sent any time one of the <code>.config</code> files is changed and you want the system to read the changes. |
| TERM | Is sent any time you want to kill <code>nsd</code> . You must send a catchable signal. SIG TERM nsd forces it to die. |

Do not use `kill -kill`, `kill -9`, or `kill -all` as it may prevent you from unmounting the file system.

Verifying `/ns/.local`

To check that the client files are installed, give this command:

```
# cat /ns/.local/hosts.byname/localhost
```

You should find output similar to the following:

```
127.0.0.1 local host
```


UUCP

UUCP stands for “UNIX to UNIX Copy Program.” It is a set of utilities that lets computers using versions of the UNIX operating system (such as IRIX) communicate with each other and with remote terminals. These utilities range from those used to copy files between computers (*uucp* and *uuto*), to those used for simple encoding and decoding (*uuencode* and *uudecode*), to those used for remote login and command execution (*cu* and *uux*).

The following topics are covered in this chapter:

- “Choosing TCP/IP or UUCP” on page 206
- “Hardware Requirements for UUCP” on page 207
- “UUCP Commands” on page 207
- “UUCP Daemons” on page 209
- “Supporting UUCP Databases” on page 210
- “UUCP Administrative Files” on page 233
- “Setting Up UUCP” on page 235
- “UUCP Error Messages” on page 247

The UUCP system is contained in the *eo*e subsystem of your IRIX distribution, in the package called *eo*e.*sw*.*uucp*. Use the *versions* command to determine if you have this subsystem installed.

UUCP connections using telephone lines and modems are used to distribute electronic mail and “net news” among thousands of computers in the USENET network.

As an administrator, you need to be familiar with the administrative tools, logs, and database files used by UUCP. This chapter provides details about the UUCP files, directories, daemons, and commands.

Choosing TCP/IP or UUCP

This section compares UUCP and the TCP/IP protocol suite for various purposes. You can use them together, each for the tasks for which it is best suited. Both UUCP and TCP/IP software are standard features of the IRIX operating system. To use the TCP/IP software, you must have one of these communications mechanisms:

- a connection to an Ethernet network
- the optional FDDI hardware and software
- the Serial Line Internet Protocol (SLIP) software

To use UUCP, you must be connected to a serial network or to any TCP/IP network.

TCP/IP provides reliable interactive and batch services. UUCP is a batch-mode service; when you enter a *uucp* command, it is placed in a queue with other commands. The system checks the queue at regular intervals and executes the commands that it finds. After your command is carried out, UUCP reports the results of the command. The time it takes to carry out a command on a remote station varies on different stations. Table 8-1 shows a comparison of features of TCP/IP and UUCP.

Table 8-1 Comparison of TCP/IP and UUCP

TCP/IP Features	UUCP Features
Runs on Ethernet and FDDI, and over serial lines	Runs over serial lines or over TCP/IP links
Transfers files interactively	Transfers files in batch mode
Executes commands on remote stations interactively	Executes commands on remote stations in batch mode
Sends mail interactively or in batch mode	Sends mail in batch mode
Starts a shell on a remote station	Starts a shell on a remote station
Provides remote login facilities with <i>rlogin/telnet</i>	Provides remote login facilities with <i>cu</i>
Transfers data to any station running TCP/IP	Transfers data to any station running UUCP

Hardware Requirements for UUCP

Before your computer can communicate with other computers through UUCP, you must set up the hardware to complete the communications link. The cables and other hardware you need depend on how you want to connect the computers: direct links, telephone lines, or local area networks.

Note: Refer to *IRIX Admin: Peripheral Devices* and the *Personal System Administration Guide* for information on setting up modems and other hardware.

Direct links You can create a direct link to another computer by running cables between serial ports on the two computers. Direct links are useful if 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. Transfer rates of up to 38,400 bits per second (bps) are possible when computers are directly linked. Such direct links are now rarely used because local area networks provide faster, easier-to-use connections.

Telephone lines Using a modem capable of dialing telephone numbers, your computer can communicate with other computers over standard phone lines. The modem dials the telephone number requested by the networking utilities. The computer it is trying to contact must have a modem capable of answering incoming calls.

UUCP Commands

UUCP programs can be divided into two categories: user programs and administrative programs. The subsections that follow describe the programs in each category.

UUCP User Program Commands

The UUCP user programs are in `/usr/bin`. No special permission is needed to use these programs; they are all described in the reference pages.

<i>cu</i>	Connects your computer to a remote computer so you can log in to that computer, allowing you to transfer some files or execute commands on either computer without dropping the initial link.
<i>uucp</i>	Lets you copy a file from one computer to another. This program creates work files and data files, queues the job for transfer, and calls the <i>uucico</i> daemon, which in turn attempts to contact the remote computer.
<i>uuto</i>	Copies files from one computer to a public spool directory on another computer (<i>/var/spool/uucppublic/receive</i>). Unlike <i>uucp</i> , which lets you copy a file to any accessible directory on the remote computer, <i>uuto</i> places the file in an appropriate spool directory and sends mail to the remote user who requested that it be picked up with <i>uupick</i> .
<i>uupick</i>	Retrieves the files placed under <i>/var/spool/uucppublic/receive</i> when files are transferred to a computer that is using <i>uuto</i> .
<i>uux</i>	Creates the work, data, and execute files needed to execute commands on a remote computer. The work file contains the same information as work files created by <i>uucp</i> and <i>uuto</i> . The execute files contain the command string to be executed on the remote computer and a list of the data files. The data files are those files required for the command's execution.
<i>uustat</i>	Displays the status of requested transfers (<i>uucp</i> , <i>uuto</i> , or <i>uux</i>). This program also provides a way to control queued transfers.

UUCP Administrative Program Commands

Most of the administrative programs are in */usr/lib/uucp*, though the UUCP database files reside in */etc/uucp*. The only exception is *uulog*, which is in */usr/bin*. These commands are described in their respective reference pages.

You should use the *uucp* login ID when you administer UUCP because it owns the basic networking and spooled data files. The home directory of the *uucp* login ID is */usr/lib/uucp*. The other UUCP login ID is *nuucp*, used by remote computers that do not have their own login IDs to access your computer. A computer that logs in with *nuucp* receives *uucico* as its shell.

The following programs are the administrative utilities of UUCP:

<i>uulog</i>	Displays the contents of a specified computer's log files. A log file is created for each remote computer with which your computer communicates. The log files contain records of each use of <i>uucp</i> , <i>uuto</i> , and <i>uux</i> .
<i>uucleanup</i>	Cleans up the spool directory. This command is normally executed from a shell script called <i>uudemon.cleanup</i> , which is started by <i>cron</i> .
<i>Uutry</i>	Tests call-processing capabilities and does a moderate amount of debugging. This command invokes the <i>uucico</i> daemon, in debug mode, to establish a communications link between your computer and the remote computer that you specify.
<i>uuccheck</i>	Checks for the presence of UUCP directories, programs, and support files. This program can also check certain parts of the <i>Permissions</i> file for obvious syntactic errors.
<i>genperm</i>	Generates the <i>Permissions</i> file for stations that assign each remote station its own login ID.

The following programs are used for initializing different types of modems. The use of these programs is described in *IRIX Admin: Peripheral Devices*, Chapter 1, "Terminals and Modems."

<i>fix-dsi</i>	This program initializes DSI modems
<i>fix-hayes</i>	This program initializes Hayes modems.
<i>fix-intel</i>	This program initializes Intel modems.
<i>fix-telebit</i>	This program initializes Telebit modems.
<i>fix-usr</i>	This program initializes U. S. Robotics modems.
<i>fix-zyxel</i>	This program initializes Telebit modems.

UUCP Daemons

There are several daemons in UUCP. These daemons handle file transfers and command executions. They can also be run manually from the shell.

<i>uucico</i>	Selects the device used for the link, establishes the link to the remote computer, performs the required login sequence and permission checks, transfers data and execute files, logs results, and notifies the user by
---------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

mail of transfer completions. It also starts *uuxqt* to execute any requested commands. When the local *uucico* daemon calls a remote computer, it “talks” to the *uucico* daemon on the remote computer during the session.

The *uucico* daemon is executed by the *uucp*, *uuto*, and *uux* programs, after all the required files have been created, to contact the remote computer. It is also executed by the *uusched* and *Uutry* programs.

<i>uuxqt</i>	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, <i>uuxqt</i> opens it to get the list of data files that are required for the execution. It then checks to see if the required data files are available and accessible. If the files are present and can be accessed, <i>uuxqt</i> checks the <i>Permissions</i> file to verify that it has permission to execute the requested command. The <i>uuxqt</i> daemon is executed by the <i>uudemon.hour</i> shell script, which is started by <i>cron</i> .
<i>uusched</i>	Schedules the queued work in the spool directory. Before starting the <i>uucico</i> daemon, <i>uusched</i> randomizes the order in which remote computers are called. <i>uusched</i> is executed by a shell script called <i>uudemon.hour</i> , which is started by <i>cron</i> .
<i>uugetty</i>	This program is very similar to the <i>getty</i> program, except that it permits a line (port) to be used in both directions. <i>uugetty</i> is assigned to a port in the <i>/etc/inittab</i> file if you want a port to be bi-directional. <i>uugetty</i> is described on the <i>uugetty(1M)</i> reference page.

Supporting UUCP Databases

The UUCP support files are in the */etc/uucp* directory.

<i>Devices</i>	Contains information concerning the location and line speed of the automatic call units (modems) and direct links. Details are in “UUCP Devices File” on page 211.
<i>Dialers</i>	Contains character strings required to negotiate with automatic call units (ACUs) or modems in establishing connections to remote computers. Details are in “UUCP Dialers File” on page 216.

<i>Systems</i>	Contains information needed by the <i>uucico</i> daemon and the <i>cu</i> program to establish a link to a remote computer. This file contains information such as the name of the remote computer, the name of the connecting device associated with the remote computer, when the computer can be reached, the telephone number, the login ID, and the password. Details are in “UUCP Systems File” on page 219.
<i>Dialcodes</i>	Contains dial-code abbreviations that can be used in the phone number field of Systems file entries. Details are in “UUCP Dialcodes File” on page 223.
<i>Permissions</i>	Defines the level of access that is granted to remote users using <i>uucp</i> or <i>uux</i> when they attempt to transfer files or remotely execute commands on your computer. Details are in “UUCP Permissions File” on page 223.
<i>Poll</i>	Defines computers that are to be polled by your station and when they are polled. Details are in “UUCP Poll File” on page 231.
<i>Sysfiles</i>	Assigns different or multiple files to be used by <i>uucico</i> and <i>cu</i> , such as <i>Systems</i> , <i>Devices</i> , and <i>Dialers</i> files. Details are in “UUCP Sysfiles File” on page 231.

There are several other files that can be considered part of the supporting database; these files are not directly related to the process of establishing a link and transferring files. The files, *Maxuuxqts*, *Maxuuscheds*, and *remote.unknown*, are described briefly in “Other UUCP Files” on page 232.

UUCP Devices File

The *Devices* file (*/etc/uucp/Devices*) contains information for all the devices that can be used to establish a link to a remote computer, such as automatic call units, direct links, and network connections.

Note: This file works interdependently with the *Dialers*, *Systems*, and *Dialcodes* files. Before you make changes in any of these files, you should be familiar with them all. A change to an entry in one file may require a change to a related entry in another file.

Each entry in the *Devices* file has the following format:

```
Type Line Line2 Class Dialer-Token-Pairs
```

Entries for use with modems should always have the form

```
Name device null speed 212 x dialer
```

Entries for use over TCP/IP network connections have the form

```
TCP - - Any TCP uucp
```

Devices file fields are defined in the following sections:

- “UUCP Type Field” on page 212
- “UUCP Line Field” on page 213
- “UUCP Line2 Field” on page 213
- “UUCP Class Field” on page 213
- “UUCP Dialer-Token-Pairs Field” on page 214
- “UUCP Device Protocols” on page 215.

UUCP Type Field

The keyword used in the Type field is matched against the third field of *Systems* file entries. The Type field can contain one of these keywords: **Direct**, **ACU**, or a **system name**.

<i>Direct</i>	This keyword indicates a direct link to another computer or a switch (for <i>cu</i> connections only).
<i>ACU</i>	This keyword indicates that the link to a remote computer is made through an automatic call unit (automatic-dial modem).
<i>Sys-Name</i>	This value indicates a direct link to a particular computer. (Sys-Name is replaced by the name of the computer.) This naming scheme is used to convey the fact that the line associated with this <i>Devices</i> entry is for a particular computer in the <i>Systems</i> file.

You can designate a protocol to use for a device within this field. See the “Protocols” section at the end of the description of this file.

UUCP Line Field

This field contains the device name of the line (port) associated with the *Devices* entry. For instance, if the automatic dial modem for a particular entry is attached to the */dev/ttyf5* line, the name entered in this field is *ttf5*.

You should always use the *ttf* devices when working with modems. These devices support hardware flow control, which is used by all modems that support V.32 or V.32bis.

UUCP Line2 Field

If the keyword **ACU** is used in the *Type* field and the ACU is an 801-type dialer, the *Line2* field contains the device name of the 801 dialer. (801-type ACUs do not contain a modem. Therefore, a separate modem is required and must be connected to a different line, defined in the *Line* field.) The need for a separate modem line means that one line would be allocated to the modem and another to the dialer. Because non-801 dialers do not normally use this configuration, they ignore the *Line2* field, but it must still contain a hyphen (-) or the word **null** as a place holder. A place holder is necessary for most modems.

UUCP Class Field

The keyword used in the *Class* field of the *Devices* file is matched against the fourth field of *Systems* file entries:

```
Devices: ACU ttyf5 null D9600 212 x telebit
```

```
Systems: eagle Any ACU D9600 14155551212 login:nuucp password:Oakgrass
```

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 will match any speed requested in a *Systems* file entry. If this field is **Any** and the *Systems* file *Class* field is **Any**, the speed defaults to 9600 bps. If the keyword **ACU** or **Direct** is used in the *Type* field, the *Class* field might contain only the speed of the device. However, the speed can also be preceded by a letter (for example, *C9600*, *D9600*) to differentiate between classes of dialers (Centrex or Dimension PBX). Including the dialer class is necessary in larger offices that have more than one type of telephone network: One network may be dedicated to serving only internal communications while another handles external. In such a case, it becomes necessary to distinguish which line(s) should be used for internal and which for external.

UUCP Dialer-Token-Pairs Field

This field contains pairs of dialers and tokens. The Dialer portion may be the name of an automatic-dial modem, or "Direct" for a direct-link device. You can have any number of Dialer-Token-Pair (DTP) fields. The Token portion may be supplied immediately following the Dialer portion; if not present, the Token portion will be taken from a related entry in the Systems file.

This field has the format

```
dialer token dialer token
```

The last pair may or may not be present, depending on the associated device (dialer). In most cases, the last pair contains only a Dialer portion and the Token portion is retrieved from the Phone field of the *Systems* file entry. A valid entry in the Dialer portion may be defined in the *Dialers* file.

The DTP field can be structured in different ways, depending on the device associated with the entry.

If an automatic-dial modem is connected directly to a port on your computer, the DTP field of the associated *Devices* file entry will have only one pair. This pair is normally 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:

```
Devices: ACU ttyf2 null 9600 212 x telebit
```

```
Dialers: telebit =&-% " " \r\p\r\c $ <K\T%\r>\c ONLINE!
```

Notice that only the Dialer portion (*telebit*) 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 case the phone number) is taken from the Phone field of a *Systems* file entry. (\T is implied)

If a direct link is established to a particular computer, the DTP field of the associated entry contains the keyword **Direct**. This is true for both types of direct-link entries, *Direct* and *System-Name*.

If an automatic-dial modem is connected to a switch, your computer must first access the switch; the switch then 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:

```

Devices: ACU ttyf2 null 9600 212 x t2500 telebit T25
Dialers: telebit "" "" \pr\ps\c est:\007 \E\D\e \007
Dialers: T25 =&-% "" \r\p\r\c $ <K\T%\r>\c ONLINE!

```

In the first pair, *t2500* is the *Dialer* and *telebit* is the *token* that is passed to the Develcon switch to tell it which device (telebit modem) to connect to your computer. This token is unique for each modem switch since each switch may be set up differently. Once the telebit modem has been connected, the second pair is accessed, where *T25* is the dialer and the token (the telephone number) is retrieved from the *Systems* file. (See the discussion of the *Systems* file's Phone field in "UUCP Systems File" on page 219.)

Two escape characters can appear in a DTP field:

- \T Indicates that the Phone (Token) field should be translated by means of the *Dialcodes* file.
- \D Indicates that the Phone (Token) field should not be translated by means of the *Dialcodes* file. If no escape character is specified at the end of a Devices entry, the \D is assumed (default).

UUCP Device Protocols

You can define the protocol to use with each device. In most cases it is not needed because you can use the default or define the protocol with the particular station you are calling. (See the discussion of the *Systems* file, specifically the Type field.) If you do specify the protocol, you must do it in the form Type, Protocol. Available protocols are

- g** This protocol is slower and more reliable than **e**. It is good for transmission over noisy telephone lines. This is the default protocol.
- e** This protocol is faster than **g**, but it assumes error-free transmission, such as over a TCP/IP network connection.
- t** This protocol, like **e**, is for use in an error-free environment, such as a TCP/IP network connection. The **t** protocol is used by systems running BSD UNIX operating systems.

UUCP Dialers File

The *Dialers* file (*/etc/uucp/Dialers*) specifies the initial conversation that must take place on a line before it can be made available for transferring data. This conversation is usually a sequence of ASCII strings that is transmitted and expected, and it is often used to dial a phone number with an ASCII dialer (such as an automatic-dial modem).

As shown in the preceding section, the fifth field in a *Devices* file entry is an index into the *Dialers* file or a special dialer type. An attempt is made 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 negotiations.

Each entry in the *Dialers* file has the following format:

```
dialer substitutions expect-send ...
```

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 technique is usually used to translate the equal sign (=) and the hyphen (-) characters into whatever the dialer requires for “wait for dial tone” and “pause.”

The *expect-send* fields are character strings.

Details of the UUCP *Dialers* file are in the next subsections:

- “UUCP Dialers File Escape Characters” on page 217
- “UUCP Dialers File Example” on page 218

UUCP Dialers File Escape Characters

The list in Table 8-2 describes some of the escape characters used in the *Dialers* file and the escape sequences used by UUCP.

Table 8-2 UUCP Escape Sequences

Escape Sequence	Meaning
\N	Send or expect a null character (ASCII NUL).
\b	Send or expect a backspace character.
\c	If at the end of a string, suppress the newline that is normally sent. Ignored otherwise.
\d	Delay two seconds before sending or reading more characters.
\p	Pause for approximately .25 to .5 seconds.
\E	Start echo checking. (From this point on, whenever a character is transmitted, login processing will wait for the character to be received before proceeding.)
\e	Turn off echo checking.
\n	Send a newline character.
\r	Send or expect a carriage return.
\s	Send or expect a space character.
\t	Send or expect a tab character.
\\	Send or expect a backslash (\) character.
EOT	Send or expect EOT newline twice.
BREAK	Send or expect a break character.
\K	Same as BREAK.
\ddd	Send or expect the character represented by the octal digits <i>ddd</i> .

UUCP Dialers File Example

Here is a sample line from 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
```

The *penril* entry in the *Dialers* file is executed as follows. The phone number argument is translated, replacing any equal sign with a **W** (wait for dial tone) and replacing any hyphen with a **P** (pause).

The handshake given by the remainder of the line works as follows:

""	Wait for nothing; in other words, proceed to the next step.
\d	Delay for two seconds.
>	Wait for a “greater than” sign (>).
Q\c	Send a Q with no terminating new line.
:	Wait for a colon (:).
\d-	Delay for two seconds.
>	Wait for a “greater than” sign (>).
s\p9\c	Send an s, pause for 0.5 second, send a 9, send no terminating new line.
)-W\p\r\ds\p9\c-)	Wait for a closing parenthesis [)]; if it is not received, process the string between the hyphens as follows: Send a W, pause, send a carriage return, delay, send an s, pause, send a 9 without a new line, and then wait for the closing parenthesis.
y\c	Send a y without a new line.
:	Wait for a colon (:).
\E\TP	Enable echo checking. (From this point on, whenever a character is transmitted, handshake processing waits for the character to be received before proceeding.) Then send the phone number. The \T instructs the program to take the phone number passed as an argument, apply the <i>Dialcodes</i> translation and the modem function translation specified by field two of this entry, then send a P .
9\c	Send a 9 without a new line.
OK	Wait for the string OK .

UUCP Systems File

The *Systems* file (*/etc/uucp/Systems*) contains the information needed by the *uucico* daemon to establish a communications link to a remote computer. Each entry in the file represents a computer that can call or be called by your computer. In addition, UUCP software by default is configured to prevent any computer that does not appear in this file from logging in to your computer. (Refer to “Other UUCP Files” on page 232 for a description of the *remote.unknown* file.) More than one entry may be present for a particular computer. The additional entries represent alternative communications paths that are tried in sequential order.

Using the *Sysfiles* file, you can define several files to be used as *Systems* files. See “UUCP Sysfiles File” on page 231 for details.

Each entry in the *Systems* file has the following format:

```
System-name Time Type Class Phone Login
```

These fields are defined in the following sections:

- “UUCP Systems File System-Name Field” on page 219
- “UUCP Systems File Time Field” on page 219
- “UUCP Systems File Type Field” on page 220
- “UUCP Systems File Class Field” on page 221
- “UUCP Systems File Phone Field” on page 221
- “UUCP Systems File Login Field” on page 222

UUCP Systems File System-Name Field

This field contains the node name of the remote computer.

UUCP Systems File Time Field

This field is a string that indicates the day-of-week and time-of-day when the remote computer can be called. The format of the Time field is

```
day[ time ] [ ; retry ]
```

The *day* portion is a list of one or more day specifiers. These specifiers are listed below.

Su, Mo, Tu, We, Th, Fr, Sa

These abbreviations stand for individual days of the week: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday.

Wk Specifies any weekday, Monday through Friday.

Any Specifies any day.

Never Specifies that this station never initiates calls to the remote computer. If the Time field is **Never**, your computer will never initiate a call to the remote computer. The call must be initiated by the remote computer. In other words, your computer is in a passive mode with respect to the remote computer. (For more information on permissions, see “UUCP Permissions File” on page 223.)

Here is an example of a Time field:

```
Wk1700-0800,Sa,Su
```

This example allows calls from 5:00 p.m. to 8:00 a.m., Monday through Friday, and any time Saturday and Sunday. The example would be an effective way to call only when phone rates are low, if immediate transfer is not critical.

The *time* portion should be a range of times such as 0800–1230. 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 at times between 6 a.m. and 8 a.m.

An optional subfield, *retry*, is available to specify the minimum time (in minutes) before a retry, following a failed attempt. The default wait is 5 minutes after the first failure, 10 after the second, and so on until a delay of about 24 hours. If the retry subfield is present, that wait is used after every failure. 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.”

UUCP Systems File Type Field

This field contains the device type that should be used to establish the communications link to the remote computer. The keyword used in this field is matched against the first field of *Devices* file entries:


```
Systems: eagle Any ACU,g D1200 3251 login:nuucp password: Oakgrass
Devices: ACU ttym2 - D1200 penril
```

You can define the protocol used to contact the station by adding it on to the Type field. The example just given shows how to attach the protocol **g** to the device type **ACU**. For direct connects, use the name of the station to which you are connecting. See “UUCP Device Protocols” on page 215.

UUCP Systems File Class Field

This field is used to indicate the transfer speed of the device used to establish the communications link. It may contain a letter and speed (for example, *C1200*, *D1200*) to differentiate between classes of dialers. (See the discussion of the Class field in “UUCP Devices File” on page 211.) Some devices can be used at any speed, so the keyword *Any* may be used. This field must match the Class field in the associated *Devices* file entry as shown here:

```
Systems: eagle Any ACU D1200 NY3251 login:nuucp password:Oakgrass
Devices: ACU ttym2 - D1200 penril
```

If information is not required for this field, use a hyphen as a place holder for the field.

UUCP Systems File Phone Field

This field is used to provide the phone number (token) of the remote computer for automatic dialers. The phone number is made up 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. For example:

```
Systems: eagle Any ACU D1200 NY3251 login:nuucp password: Oakgrass
Dialcodes: NY 9=1212555
```

In this string, an equal sign (=) tells the ACU to wait for a secondary dial tone before dialing the remaining digits. A hyphen (-) in the string instructs the ACU to pause four seconds before dialing the next digit.

If your computer is connected to a modem switch, you may access other computers that are connected to that switch. The *Systems* file entries for these computers does not have a phone number in the Phone field. Instead, this field contains the token that must be passed on to the switch so it will know which computer your computer wishes to communicate with. This token is usually just the station 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 means of the *Dialcodes* file.

UUCP Systems File Login Field

This field contains login information given as a series of fields and subfields of this format:

expect send

The *send* string is sent when the *expect* string is received. The expect field may be made up of subfields of this form:

expect[-send-expect] . . .

The send field is sent if the prior expect is not successfully read and the expect following the send is the next expected string. For example, with login--login, UUCP expects the string login. If UUCP receives that string, it goes on to the next field. If it does not receive the string, it sends nothing, followed by a new line, then looks for login again. If no characters are initially expected from the remote computer, the characters "" (null string) should be used in the first expect field. Note that all send fields are sent followed by a newline unless the send string is terminated with a `\c`.

Here is an example of a *Systems* file entry that uses an expect-send string:

```
owl Any ACU 1200 NY6013 "" \r login:-BREAK-login: uucpx word: xyzy
```

This example means causes UUCP to send a carriage return and wait for the string login:. If it doesn't receive the string, it sends a break character and waits for login: again. When it does receive the string login, it sends the login name **uucpx**; waits to receive the string word: (the last part of "Password:"); and sends the password **xyzy**.

Several escape sequences cause specific actions when they are a part of a string sent during the login sequence. These are the same escape sequences used in the *Dialers* file, listed in Table 8-2.

UUCP Dialcodes File

The *Dialcodes* file (*/etc/uucp/Dialcodes*) contains the dial-code abbreviations that can be used in the Phone field of the *Systems* file. Each entry has the following format:

```
abb dial-seq
```

abb is the abbreviation used in the *Systems* file Phone field and *dial-seq* is the dial sequence that is passed to the dialer when that *Systems* file entry is accessed.

For example, the following entry would work with a Phone field in the *Systems* file such as *jt7867*:

```
jt 9=555-
```

When the entry containing *jt7867* was encountered, the sequence *9=555-7867* would be sent to the dialer if the token in the dialer-token-pair was *\T*.

UUCP Permissions File

The *Permissions* file (*/etc/uucp/Permissions*) specifies the permissions that remote computers have with respect to login, file access, and command execution. There are options that restrict the remote computer's ability to request files and its ability to receive files queued by the local site. Another option specifies the commands that a remote site can execute on the local computer.

The program */etc/uucp/genperm* is recommended for creating a sample or default *Permissions* file from the *Systems* file.

These subsections explain details of the *Permissions* file:

- “UUCP Permissions File Entry Structure” on page 224
- “UUCP Permissions File Considerations” on page 224
- “UUCP Permissions File Options” on page 225

UUCP Permissions File Entry Structure

Each entry is a logical line with physical lines terminated by a backslash (\) to indicate continuation. (Note that such continuations are not possible in most other UUCP files.) Entries are made up of options delimited by white space. Each option is a name/value pair in the following format:

name=value

Note that no white space is allowed within an option assignment.

Comment lines begin with a number sign (#) and occupy the entire line up to a newline character. Blank lines are ignored (even within multi-line entries).

There are two types of *Permissions* file entries:

- | | |
|---------|-----------------------------------------------------------------------------------------------------|
| LOGNAME | Specifies the permissions that take effect when a remote computer logs in to (calls) your computer. |
| MACHINE | Specifies the permissions that take effect when your computer logs in to (calls) a remote computer. |

LOGNAME entries begin with a LOGNAME option and MACHINE entries begin with a MACHINE option.

UUCP Permissions File Considerations

Keep these rules in mind when using the *Permissions* file to restrict the level of access granted to remote computers:

- Any login ID used by a remote computer to log in for UUCP communications must appear in one and only one LOGNAME entry.
- Any site that is called whose name does not appear in a MACHINE entry will have the following default permissions and restrictions:
 - Local send and receive requests are executed.
 - The remote computer can send files to your computer's */var/spool/uucppublic* directory.
 - The command sent by the remote computer for execution on your computer must be one of the default commands, usually *rmail*.

UUCP Permissions File Options

This section describes each option, specifies how it is used, and lists its default value.

REQUEST 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 that follows specifies that the remote computer can request to transfer files from your computer:

```
REQUEST=yes
```

The following string specifies that the remote computer cannot request to receive files from your computer:

```
REQUEST=no
```

This is the default value. It will be used if the **REQUEST** option is not specified. The **REQUEST** option can appear in either a **LOGNAME** (remote calls you) entry or a **MACHINE** (you call remote) entry.

A note on security: When a remote computer calls you, you cannot verify its identity unless you have a unique login and password for that computer.

SENDFILES When a remote computer calls your computer and completes its work, it may attempt to take work your computer has queued for it. The **SENDFILES** option specifies whether your computer can send the work queued for the remote computer.

The string shown here specifies that your computer may send the work that is queued for the remote computer as long as it logged in as one of the names in the **LOGNAME** option:

```
SENDFILES=yes
```

This string is mandatory if your computer is in a “passive mode” with respect to the remote computer.

The string that follows specifies that files queued in your computer will be sent only when your computer calls the remote computer:

```
SENDFILES=call
```

The call value is the default for the **SENDFILE** option. This option is significant only in **LOGNAME** entries, because **MACHINE** entries apply when calls are made to remote computers. If the option is used with a **MACHINE** entry, it will be ignored.

READ and **WRITE**

These options specify the various parts of the filesystem that *uucico* can read from or write to. The **READ** and **WRITE** options can be used 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
```

These strings specify permission to access any file that can be accessed by a local user with “other” permissions:

```
READ=/ WRITE=/
```

Because this suggestion may compromise security, use it only if required.

The value of these entries is a colon-separated list of pathnames. The **READ** option is for requesting files, and the **WRITE** option for depositing files. One of the values must be the prefix of any full pathname of a file coming in or going out. To grant permission to deposit files in */usr/news* as well as in the public directory, the following values would be used with the **WRITE** option:

```
WRITE=/var/spool/uucppublic:/usr/news
```

Note that if you use the **READ** and **WRITE** options, you must specify all pathnames because the pathnames are not added to the default list. For instance, if the */usr/news* pathname were the only one specified in a **WRITE** option, permission to deposit files in the public directory would be denied.

You should be careful which directories you make accessible for reading and writing by remote stations. For example, you probably wouldn't want remote computers to be able to write over your */etc/passwd* file, so */etc* shouldn't be open to writes.

NOREAD and NOWRITE

The **NOREAD** and **NOWRITE** options specify exceptions to the **READ** and **WRITE** options or defaults. The strings shown here would permit one remote computer to read any file except those in the */etc* directory (and its subdirectories—remember, these are prefixes) and to write only to the default */var/spool/uucppublic* directory:

```
READ=/ NOR EAD=/etc WRITE=/var/spool/uucppublic
```

NOWRITE works in the same manner as the **NOREAD** option. **NOREAD** and **NOWRITE** can be used in both **LOGNAME** and **MACHINE** entries.

CALLBACK

The **CALLBACK** option is used in **LOGNAME** entries to specify that no transaction will take place until the calling station is called back. You would use **CALLBACK** for two reasons: From a security standpoint, if you call back a station you can be sure it is the station it says it is. If you are doing long data transmissions, you can choose the station that will be billed for the longer call.

The string that follows specifies that your computer must call the remote computer back before any file transfers will take place:

```
CALLBACK=yes
```

The default for the **CALLBACK** option is

```
CALLBACK=no
```

The **CALLBACK** option is very rarely used. Note that if two sites have this option set for each other, a conversation cannot be started.

COMMANDS The **COMMANDS** option can be hazardous to the security of your station. Use it with extreme care.

The *uux* program generates remote execution requests and queues them to be transferred to the remote computer. Files and a command are sent to the target computer for remote execution.

The **COMMANDS** option can be used in **MACHINE** entries to specify the commands that a remote computer can execute on your computer. Note that **COMMANDS** is not used in a **LOGNAME** entry; **COMMANDS** in **MACHINE** entries defines command permissions, whether you call the remote station or it calls you.

This string indicates the default commands that a remote computer can execute on your computer:

```
COMMANDS=rmail
```

If a command string is used in a `MACHINE` entry, the default commands are overridden. For instance, in the following example, the entry overrides the **COMMANDS** default so that the computers *eagle*, *owl*, and *hawk* can now execute *rmail* and *rnews* on your computer:

```
MACHINE=eagle:owl:hawk REQUEST=yes  
COMMANDS=rmail:/usr/bin/rnews  
READ=/ WRITE=/
```

In addition to the names as specified above, there can be full pathnames of commands. For example, this line specifies that command *rmail* use the default path:

```
COMMANDS=rmail:/usr/bin/rnews:/usr/local/lp
```

The default paths for your computer are */bin*, */usr/sbin*, */usr/bsd*, and */usr/bin*. When the remote computer specifies *rnews* or */usr/bin/rnews* for the command to be executed, */usr/bin/rnews* is executed, regardless of the default path. Likewise, */usr/local/lp* is the *lp* command that is executed.

Note: Including the `ALL` value in the list means that any command from the remote computer(s) specified in the entry is executed. If you use this value, you give the remote computer full access to your computer. *Be careful.* This value allows far more access than normal users have.

This string illustrates the greater access:

```
COMMANDS=/usr/bin/rnews:ALL:/usr/local/lp
```

Two points about this string should be noted. The `ALL` value can appear anywhere in the string, and the pathnames specified for *rnews* and *lp* will be used (instead of the default) if the requested command does not contain the full pathnames for *rnews* or *lp*.

The **VALIDATE** option should be used with the **COMMANDS** option whenever potentially dangerous commands like *cat* and *uucp* are specified with the **COMMANDS** option. Any command that reads or writes files is potentially dangerous to local security when executed by the UUCP remote execution daemon (*uuxqt*).

VALIDATE The **VALIDATE** option is used in conjunction with the **COMMANDS** option when specifying commands that are potentially dangerous to your computer's security. It is used to provide a certain degree of verification of the caller's identity. The use of the **VALIDATE** option requires that privileged computers have a unique login and password for UUCP transactions. An important aspect of this validation is that the login and password associated with this entry be protected. If an outsider gets that information, that particular **VALIDATE** option can no longer be considered secure. (**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**.)

Give careful consideration to providing a remote system with a privileged login and password for UUCP transactions. Giving another system these privileges is like giving anyone on that computer a normal login and password on your computer. Therefore, if you cannot trust everyone at the remote site, do not provide that system with a privileged login and password.

LOGNAME The **LOGNAME** option ensures that remote stations attempting to log in to your computer have login privileges. The following **LOGNAME** entry specifies that if one of the remote computers that claims to be *eagle*, *owl*, or *hawk* logs in to your computer, it must have used the login *uucpfriend*:

```
LOGNAME=uucpfriend VALIDATE=eagle:owl:hawk
```

As can be seen, if an outsider gets the *uucpfriend* login and password, marauding is trivial.

But what does this have to do with the **COMMANDS** option, which appears only in **MACHINE** entries? It links the **MACHINE** entry (and **COMMANDS** option) with a **LOGNAME** entry associated with a privileged login. This link is needed because the execution daemon is not running while the remote computer is logged in. In fact, it is an asynchronous process with no knowledge of 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 computer. These spool directories have write permission 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.

The following example shows the relationship between the **MACHINE** and **LOGNAME** entries:

```
MACHINE=eagle:owl:hawk REQUEST=yes \
COMMANDS=rmail:/usr/bin/rnews \
READ=/ WRITE=/
LOGNAME=uucpz VALIDATE=eagle:owl:hawk \
REQUEST=yes SENDFILES=yes \
READ=/ WRITE=/
```

The value in the **COMMANDS** option means that remote mail and */usr/bin/rnews* can be executed by remote users.

In the first entry, you must make the assumption that when you want to call one of the computers listed, you are really calling either *eagle*, *owl*, or *hawk*. Therefore, any files 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 will also be put in the privileged spool directory. You therefore have to validate that the computer has the privileged login *uucpz*.

MACHINE Entry for “Other” Systems

You may want to specify different option values for computers your computer calls that are not mentioned in specific **MACHINE** entries. This situation may occur when there are many computers calling in, and the command set changes from time to time. The name **OTHER** for the computer name is used for this entry:

```
MACHINE=OTHER \
COMMANDS=rmail:rnews:/usr/bin/Photo:/usr/bin/xp
```

All other options available for the **MACHINE** entry may also be set for the computers that are not mentioned in other **MACHINE** entries.

Combining **MACHINE** and **LOGNAME** Entries

It is possible to combine **MACHINE** and **LOGNAME** entries into a single entry where the common options are the same. For example, the two entries that follow share the same **REQUEST**, **READ**, and **WRITE** options:

```
MACHINE=eagle:owl:hawk REQUEST=yes \
  READ=/ WRITE=/
LOGNAME=uucpz REQUEST=yes SENDFILES=yes \
  READ=/ WRITE=/
```

These two entries can be merged:

```
MACHINE=eagle:owl:hawk REQUEST=yes \
LOGNAME=uucpz SENDFILES=yes \
READ=/ WRITE=/
```

MYNAME The **MYNAME** option is used to override the name of the local computer, when the local computer identifies itself to the remote computer. This facility is useful when a computer is replaced or renamed, and its neighbors need to process old traffic to the old name.

UUCP Poll File

The *Poll* file (*/etc/uucp/Poll*) 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 (a space won't work), and finally the hours at which the computer should be called. The format of entries in the *Poll* file is

```
sys-name hour ...
```

For example, the following entry provides polling of computer *eagle* every four hours:

```
eagle 0 4 8 12 16 20
```

The *uudemon.poll* script does not actually perform the poll. It merely sets up a polling work file (always named *C.file*) in the spool directory that will be seen by the scheduler, which is started by *uudemon.hour*.

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. Here are some cases where this optional file may be useful:

- You may want to use different *Systems* files so that requests for login services can be made to phone numbers different from those used for requests for *uucp* services.

- You may want to use *Dialers* files that have different handshaking for *cu* and *uucp*.
- You may want to have multiple *Systems*, *Dialers*, and *Devices* files. The *Systems* file in particular may become large, making it more convenient to split it into several smaller files.

The format of the *Sysfiles* file is

```
service=w systems=x:x dialers=y:y devices=z:z
```

The *w* parameter is replaced by *uucico*, *cu*, or both separated by a colon; *x* is one or more files to be used as the *Systems* file, with each filename separated by a colon and read in the order presented; *y* is one or more files to be used as the *Dialers* file; and *z* is one or more files to be used as the *Devices* file. Each file is assumed to be relative to the */etc/uucp* directory, unless a full path is given. A backslash-carriage return (`\Return`) can be used to continue an entry to the next line.

Here is an example using a local *Systems* file in addition to the usual *Systems* file:

```
service=uucico:cu systems=Systems:Local_Systems
```

If this line is in */etc/uucp/Sysfiles*, then both *uucico* and *cu* will first look in */etc/uucp/Systems*. If the station they're trying to call doesn't have an entry in that file, or if the entries in the file fail, then they'll look in */etc/uucp/Local_Systems*.

When different *Systems* files are defined for *uucico* and *cu* services, your station will store two different lists of stations. You can print the *uucico* list by using the *uname* command, or the *cu* list by using the *uname -c* command.

Other UUCP Files

Three files, in addition to those described in the preceding subsections, have an impact on the use of basic networking facilities. In most cases, the default values are fine and no changes are needed. If you want to change the default values, however, use any standard IRIX text editor (*ed*, *vi*, or *jot*).

Maxuuxqts This file defines the maximum number of *uuxqt* programs that can run at once. The default number is two.

Maxuuscheds This file defines the maximum number of *uusched* programs that can run at once. The default number is two.

unknown This file is a program that executes when a station that is not in any of the *Systems* files starts a conversation. The program logs the conversation attempt and refuses the connection. If you change the permissions of this file so it cannot execute (*chmod 000 unknown*), your station will accept any conversation requests.

UUCP Administrative Files

The UUCP administrative files are created in spool directories to lock devices, hold temporary data, or keep information about remote transfers or executions.

TM (temporary data file)

These data files are created by UUCP processes under the spool directory (for example, */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 starting at 0.

When the entire file is received, the *TM.pid.ddd* file is moved to the pathname specified in the *C.sysnxxxx* file (discussed later in this section) that caused the transmission. If processing is abnormally terminated, the *TM.pid.ddd* file may remain in the *X* directory. These files should be automatically removed by *uucleanup*.

LCK (lock file) 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 names of lock files have this format:

LCK..str

str is either a device or computer name. These files may remain in the spool directory if the communications link is unexpectedly dropped (usually because of a computer crash). Lock files are ignored (removed) after the parent process is no longer active. Each lock file contains the process ID of the process that created the lock.

- C. (work file) 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 representing the grade (priority) of the work, and *xxxx* is the four-digit job sequence number assigned by UUCP. Work files contain the following information:

- Full pathname of the file to be sent or requested.
- Full pathname of the destination or user or filename.
- User login name.
- List of options.
- Name of associated data file 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.
- Login name of the remote user to be notified upon completion of the transfer.

- D. (data file) Data files are created when the command line specifies that the source file should be copied 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 and *xxxx* is a four-digit job sequence number assigned by UUCP. The four-digit job sequence number may be followed by a subsequence number, *yyy*, 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 representing the grade (priority) of the work, and *xxxx* is a four-digit sequence number assigned by UUCP. Execute files contain the following information:

- requester's login and computer name
- name of file(s) required for execution
- input to be used as the standard input to the command string
- filename for a file to receive standard output from the command execution, together with the hostname of the computer on which the file resides
- command string
- option lines for return status requests

Setting Up UUCP

Setting up UUCP involves five steps:

1. Determine the *remote* and *local* stations. See "Determining the Remote and Local Stations for a UUCP Connection" on page 235.
2. Make the physical connection. See "Making the Physical Connection for UUCP" on page 236.
3. Configure the *local* (calling) station. See "Configuring the Local Station for UUCP" on page 236.
4. Configure the *remote* (called) station. See "Configuring the Remote Station for UUCP" on page 240.
5. Test the UUCP connection. See "Testing the UUCP Connection" on page 243.

You can also use UUCP on a TCP/IP network. See "Setting Up UUCP on a TCP/IP Connection" on page 245.

Determining the Remote and Local Stations for a UUCP Connection

Typically, the *local* station is the station that initiates the UUCP connection. The *remote* station is the station that responds to UUCP connection requests. However, with the arrival of *uugetty* (a program that allows bi-directional line usage), the distinction between the local and remote station is usually only the station name.

Making the Physical Connection for UUCP

UUCP supports physical connections for TCP/IP local area network connections, direct links, or telephone lines. This example assumes a direct link. The procedure for running UUCP over telephone line or local area networks is similar, requiring minor adjustments to the various configuration files.

A direct link constitutes a connection between two Data Terminal Equipment (DTE) devices. The devices must be fooled into thinking that they are communicating with a Data Communication Equipment (DCE) device. The way to get around this is with a null modem.

The minimum pinning configuration shown in Table 8-3.

Table 8-3 Three-Wire Null-Modem Pinning Configuration

IRIS A	IRIS B
2 Transmit Data	3 Receive Data
3 Receive Data	2 Transmit Data
7 Signal Ground	7 Signal Ground

Attach the null modem cable to serial port two (*ttyf2*) on the local and remote workstations.

Note: For more information on cables and modems, see *IRIX Admin: Peripheral Devices*, Chapter 1, "Terminals and Modems."

Configuring the Local Station for UUCP

The remote station name in our example is *us*, and the local station name is *japan*. There are two steps in configuring the local station:

1. Update standard system files—for */etc/passwd*, see “Updating */etc/passwd* on the Local Station for UUCP” on page 237; for */etc/group*, see “Updating */etc/group* on the Local Station for UUCP” on page 238; for */etc/inittab*, see “Updating */etc/inittab* on the Local Station for UUCP” on page 238.
2. Modify the UUCP configuration files—for */etc/uucp/Systems*, see “Modifying */etc/uucp/Systems* on the Local Station for UUCP” on page 238; for *etc/uucp/Devices*, see “Modifying */etc/uucp/Devices* on the Local Station for UUCP” on page 239; for */etc/uucp/Dialers*, see “Modifying */etc/uucp/Dialers* on the Local Station for UUCP” on page 239; for */etc/uucp/Permissions*, see “Modifying */etc/uucp/Permissions* on the Local Station for UUCP” on page 240.

Updating */etc/passwd* on the Local Station for UUCP

To ensure proper security and access, you need to ensure that the user entries for *uucp* and *nuucp* are both present and correct. The *uucp* entry in the *passwd* file is for ownership purposes, and the *nuucp* entry is for remote UUCP access. Ensure that your password file has both entries and that they are the same as the following example. If the *uucp* and *nuucp* entries don't match the following, edit those accounts so they do match.

```
uucp:*:3:5:UUCP Owner:/usr/lib/uucp:/bin/csh
nuucp::10:10:Remote UUCP
User:/var/spool/uucppublic:/usr/lib/uucp/uucico
```

In the above example, the *passwd* entry for *nuucp* is split across two lines due to formatting constraints. In the actual file, the entry appears on a single line.

On a newly installed station, neither *uucp* nor *nuucp* has a password. It is a good idea to put an asterisk (*) in the password field for *uucp*, because no one should log in as *uucp*. You need to assign *nuucp* a valid password that matches the password you assign for *nuucp* in the *Systems* file. (See “UUCP *Systems* File” on page 219. For example, assign *nuucp* the password “secret.”)

```
New password: secret
Re-enter new password: secret
```

Updating `/etc/group` on the Local Station for UUCP

Check this file to ensure that there are valid groups for both `uucp` and `nuucp`. Compare your `uucp` and `nuucp` group entries with the following. If there is a discrepancy, correct it now.

```
uucp: :5:uucp
nuucp: :10:nuucp
```

Updating `/etc/inittab` on the Local Station for UUCP

This sample entry is for the local station. It allows calls to be initiated on port two, but does not allow incoming calls on the port. Edit your `/etc/inittab` entry for “t2” as follows:

```
t2:23:off:/usr/lib/uucp/uugetty -Nt 60 ttyf2 co_9600 # port 2
```

For complete information on the `uugetty` command, see the `uugetty(1M)` reference page. As usual, any time you make a change to the `/etc/inittab`, you must tell `init` to read the file again with the `telinit q` command. Enter the following command:

```
/etc/telinit q
```

Modifying `/etc/uucp/Systems` on the Local Station for UUCP

The `Systems` file contains information describing the station(s) that the local station knows about. Add the following line to the bottom of the file:

```
us Any systemx 9600 unused ogin:--ogin: nuucp ssword: \ secret
```

Note: The `Systems` file is read only, so if you are using `vi` you must force this change to be written out by exiting `vi` with the `:wq!` option.

The first field specifies the name of the station that can call (the remote station). The second field indicates that the specified station can call at any time. The third field tells `uucp` the name of the device to use (**systemx**). The third field must match one of the first field entries found in `/etc/uucp/Devices`. The fourth field specifies the transfer speed (9600). The fifth field is normally used for a phone number, but unused for direct links. The rest of the line handles the login sequence; it is the chat script negotiated between the local station and the remote station. This chat script is very important for a successful `uucp` connection.

Modifying /etc/uucp/Devices on the Local Station for UUCP

The *Devices* file contains information about the physical connection between the two stations. Remove the pound sign from the *systemx* device entry so it looks like the following:

```
# ---A direct connection to a system
systemx ttyf2 - Any direct
```

Note: If you have another direct connection to a station on another port, copy the *systemx* device entry and modify the port number accordingly.

The first field in the *Devices* file links the device to the *Systems* file (third field entry). The second field tells *uucp* which port to access. The third field is used with an Automatic Call Unit (ACU). Direct links use a dash in the third field. The fourth field specifies the line speed. The “Any” entry allows the speed to be determined by the */etc/inittab* file for that particular device. The fifth field contains the dialer name. It must be a valid entry in the */etc/uucp/Dialers* file.

Modifying /etc/uucp/Dialers on the Local Station for UUCP

This file contains the chat script for the *uucp* device. Because this is a direct connection, the chat script is picked up from the *Systems* file. However, there still has to be a valid dialers entry for the direct connection. Verify that the *Dialers* file has an entry for the “direct” dialer. Enter the following command:

```
grep direct /etc/uucp/Dialers
```

The system responds with

```
direct
# The following entry is for use with direct connections
uudirect "" "" \r\d in:--in:
```

Modifying `/etc/uucp/Permissions` on the Local Station for UUCP

The *Permissions* file controls remote *uucp* access with regard to remote users and stations. (See “UUCP Permissions File” on page 223 for descriptions of all of the options.) For this example, edit the *Permissions* file to look like the following:

```
#dent "@(#)uucp:Permissions2.2"
# This entry for public login.
# It provides the default permissions.
# See the Basic Networking Utilities Guide for more information.
LOGNAME=nuucp MACHINE=us READ=/var/spool/uucp/uucppublic \
WRITE=/var/spool/uucppublic REQUEST=yes SENDFILES=yes \
COMMANDS=rmail
```

Note: This entry must be interpreted as a single line, even if it exceeds more than one physical line.

This entry specifies that the user, *nuucp*, is allowed to log in from the remote station (*us*). The *nuucp* user on *us* may read any files that reside in `/var/spool/uucp/uucppublic` directory and write to the general public directory `/var/spool/uucppublic`. The users on *us* may make requests. Users on *japan* can send files.

Configuring the Remote Station for UUCP

The remote station name in our example is *japan*. The local station name in our example is *us*. There are two steps in configuring the remote station:

1. Update standard system files—For `/etc/passwd` see “Updating `/etc/passwd` on the Remote Station for UUCP” on page 241; for `/etc/group`, see “Updating `/etc/group` on the Remote Station for UUCP” on page 241; for `/etc/inittab` see “Updating `/etc/inittab` on the Remote Station for UUCP” on page 241.
2. Modify the UUCP configuration files—For `/etc/uucp/Systems` see “Modifying `/etc/uucp/Systems` on the Remote Station for UUCP” on page 242; for `/etc/uucp/Permissions` see “Modifying `/etc/uucp/Permissions` on the Remote Station for UUCP” on page 242.

Updating `/etc/passwd` on the Remote Station for UUCP

To ensure proper security and access, you need to ensure that the user entries for `uucp` and `nuucp` are both present and correct. The `uucp` entry in the `passwd` file is for ownership purposes and the `nuucp` entry is for remote UUCP access. Ensure that your password file has both entries and that they are the same as the following example. If the `uucp` and `nuucp` entries don't match the following, edit them so they do match:

```
uucp:*:3:5:UUCP Owner:/usr/lib/uucp:/bin/csh
nuucp::10:10:Remote UUCP User:/var/spool/uucppublic:/usr/lib/uucp/uucico
```

On a newly installed station, neither `uucp` nor `nuucp` has a password. It is a good idea to put an asterisk (*) in the password field for `uucp`, because no one should log in as `uucp`. You need to assign `nuucp` a valid password that matches the password you assign for `nuucp` in the Systems file. (See "UUCP Systems File" on page 219.) Assign `nuucp` the password "secret" with this command:

```
passwd nuucp
New password: secret
Re-enter new password: secret
```

Updating `/etc/group` on the Remote Station for UUCP

Check this file to ensure that there are valid groups for both `uucp` and `nuucp`. Compare your `uucp` and `nuucp` group entries with the following example. If there is a discrepancy, correct it now.

```
uucp: :5:uucp
nuucp: :10:nuucp
```

Updating `/etc/inittab` on the Remote Station for UUCP

This sample entry is for the remote station. It allows calls to be received on serial port 2, but does not allow outgoing calls on the port. Edit your `/etc/inittab` entry for "t2" as follows:

```
t2:23:respawn:/usr/lib/uucp/uugetty -Nt 60 ttyf2 co_9600#pt 2
```

For complete information on the `uugetty` command, see the `uugetty(1M)` reference page. As usual, any time you make a change to the `/etc/inittab`, you must use the `telinit q` command to tell `init` to read the file again. Enter the following command:

```
/etc/telinit q
```

Modifying /etc/uucp/Systems on the Remote Station for UUCP

The *Systems* file contains information describing the station(s) that the remote station knows about. Add this line to the bottom of the *Systems* file:

```
japan Never
```

Note: The permission of the *Systems* file is read only. If you edit this file, you may have to use a forced write in order to save your changes. Consult the *jot* or *vi* reference page for more information.

The first field specifies the name of a station that can call the local station. The second field indicates the times this station can make the call. The value **Never** in this field indicates that this station may receive calls, but never initiate them.

Modifying /etc/uucp/Permissions on the Remote Station for UUCP

The *Permissions* file controls remote *uucp* access with regard to remote users and stations. (See "UUCP Permissions File" on page 223 for descriptions of all the options.) For this example, edit the *Permissions* file to look like the following:

```
#ident"@(#)uucp:Permissions2.2"
# This entry for public login.
# It provides the default permissions.
# See the Basic Networking Utilities Guide for more
                                                                    information.
LOGNAME=nuucp MACHINE=japan READ=/var/spool/uucp/uucppublic\
WRITE=/var/spool/uucppublic REQUEST=yes SENDFILES=yes \
COMMANDS=rmail
```

Note: This entry must be interpreted as a single line, even if it expands to more than one physical line.

This entry specifies that the user, *nuucp*, is allowed to log in from the local station (*japan*). The *nuucp* user on *japan* may read any files that reside in */var/spool/uucp/uucppublic* directory and write to the general public directory */var/spool/uucppublic*. The users on *japan* may request files. The users on *us* can send files.

Testing the UUCP Connection

There are two basic tools for testing a UUCP connection:

- the *cu* program described in “Testing the UUCP Connection With *cu*” on page 243
- the *Uuetry* program described in “Testing the UUCP Connection With *Uuetry*” on page 244

Testing the UUCP Connection With *cu*

The *cu* program is used to test the basic functionality of the UUCP connection. When you use *cu* directly, you are performing the login process as if you were the *uucp* programs. The *cu* command is also used for direct modem connections for terminal emulation. The **-d** option to *cu* is used for diagnostics and causes traces of information to be printed out to the standard output (your shell window). You should always use this mode when testing a UUCP connection.

The *cu* command tests the physical connection to the remote station, UUCP configuration files, operating system files, and the *uucico* daemon on the remote station.

Note: The default permissions on devices (*/dev/ttyf2*) are set to 622. For *cu* to access your device, you need to change the permissions to 666.

1. The *cu* command must be run from the local (calling) station. Enter the *cu* command from the local station (*japan*) as follows:

```
/usr/bin/cu -d us
```

You see output similar to the following:

```
conn(us)
Device Type us wanted
mlock ttyf2 succeeded
filelock: ok
fixline(5, 9600)
processdev: calling setdevcfg(cu, us)
gdial(direct) called
getto ret 5
device status for fd=5
F_GETFL=2,iflag='12045',oflag='0',cflag='6275',lflag='0',line='1'
```

```
cc[0]='\177',[1]='\34',[2]='\10',[3]='\25',[4]='\1',[5]='\0',[6]='\0',[7]='\0',
call _mode(1)
Connected
_receive started
transmit started
```

2. When the system pauses, press **Enter**.
3. When you see the login prompt, log in as *nuucp* and supply the password for *nuucp* (in this case, the password is **secret**):

```
Break your connection with a tilde(~) dot(.) and a carriage return (<CR>).
us login: nuucp
Password: secret
IRIX Release 6.2 IP22 us
Copyright (c) 1987-1996 Silicon Graphics, Inc. All Rights Reserved.
Last login: Thu Aug 8 09:14:29 MDT 1996 on ttyf2
here=japan~[us].
call tilda(.)
```

Testing the UUCP Connection With Uutry

Uutry is the program that tests the copy-in/copy-out program (*uucico*). *uucico* must be functioning properly before you can actually transfer data. Enter the *Uutry* command from the local station (*japan*) to the remote station (*us*):

```
/usr/lib/uucp/Uutry us
```

You should see output similar to the following:

```
/usr/lib/uucp/uucico -r1 -sus -x5 >/tmp/us 2>&1&
tmp=/tmp/us
mchFind called (us)
conn(us)
Device Type us wanted
mlock ttyf2 succeeded
processdev: calling setdevcfg(uucico, us)
gdial(direct) called
getto ret 5
expect: (ogin:)
```

Note: The system may pause here for several minutes.

```
sendthem (^M)
```



```

expect: (ogin:)
^M^M^J^M^J^Jus login:got it
sendthem (nuucp^M)
expect: (ssword:)
nuucp^M^JPassword:got it
sendthem (secret^M)
Login Successful: System=us
msg-ROK
Rmtname us, Role MASTER, Ifn - 5, Loginuser - root
rmsg - 'P' got Pg
wmsg 'U'g
Proto started g
*** TOP *** - role=1, setline - X
wmsg 'H'
rmsg - 'H' got HY
PROCESS: msg - HY
HUP:
wmsg 'H'Y
cntrl - 0
send OO 0,exit code 0
Conversation Complete: Status SUCCEDED

```

When you see `Status SUCCEDED`, *Utry* has successfully tested *uucico*. Press `Ctrl+C` to break out of *Utry*.

Setting Up UUCP on a TCP/IP Connection

In many cases, you may decide to use the UUCP tools over conventional TCP/IP network connections. There are entries in the *Devices* file provided for this, but you must make some changes in the */usr/etc/inetd.conf* and */etc/uucp/Systems* files. Follow these steps:

1. Edit the */usr/etc/inetd.conf* file on the remote host and find this line:

```
#uucp stream tcp nowait root /usr/lib/uucp/uucpd uucpd
```

Remove the leading hashmark (#) to uncomment the line, and use these commands to make the change take effect:

```
/etc/init.d/network stop
/etc/init.d/network start
```

This change tells the remote system to run the *uucpd* daemon when a request comes in for UUCP transfer.

2. Add a line similar to the following to your local `/etc/uucp/Systems` file:

```
remotehost Any TCP Any
```

The name *remotehost* should be replaced with the name of the remote host you will be calling.

3. Run the `/etc/uucp/genperm` command as **root** on the local host to generate the UUCP *Permissions* file.
4. Enter the following command to check that everything is set up correctly:

```
/usr/lib/uucp/uucheck -v
```

There is a great deal of output. You should see output similar to the following for the specific entry you made:

```
When we call system(s): (remotehost)
We DO allow them to request files.
They can send files to
/var/spool/uucppublic (DEFAULT)
They can request files from
/var/spool/uucppublic
/usr/lib/mail
/usr/people/ftp
Myname for the conversation will be MyName.
PUBDIR for the conversation will be /var/spool/uucppublic.
```

```
Machine(s): (remotehost)
CAN execute the following commands:
command (rmail), fullname (/bin/rmail)
command (rnews), fullname (/usr/bin/rnews)
command (cunbatch), fullname (/usr/lib/news/cunbatch)
```

The `cu` command does not work for UUCP over a TCP connection. Use the `/usr/lib/uucp/Uutry` command instead. *Uutry* is documented completely in the *Uutry(1M)* reference page and in “Testing the UUCP Connection With *Uutry*” on page 244.

UUCP Error Messages

This section describes common error messages associated with the UUCP environment. UUCP error messages can be divided into two categories: ASSERT error messages described on “ASSERT Error Messages” on page 247; and STATUS error messages, described on “STATUS Error Messages” on page 249.

ASSERT Error Messages

When a process is aborted, the station records ASSERT error messages in */var/spool/uucp/.Admin/errors*. These messages include the filename, the *sccsid*, the *line number*, and the text listed in Table 8-4. In most cases, these errors are the result of file system problems.

Table 8-4 Assert Error Messages

Error Message	Description/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 other call failed.
CAN'T READ	A <code>read()</code> , <code>fgets()</code> , or other call failed.
CAN'T CREATE	A <code>create()</code> call failed.
CAN'T ALLOCATE	A dynamic allocation failed.
CAN'T LOCK	An attempt to make an LCK (lock) file failed. In some cases, 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.

Table 8-4 Assert Error Messages (**continued**)

Error Message	Description/Action
CAN'T MOVE TO CORRUPT DIR	An attempt to move some bad <i>C.</i> or <i>X.</i> files to the <i>/var/spool/uucp/.Corrupt</i> directory failed. The directory is probably missing or has wrong modes or owner.
CAN'T CLOSE	A close() or fclose() call failed.
FILE EXISTS	The creation of a <i>C.</i> or <i>D.</i> file was attempted, but the file already exists. This situation occurs when there is a problem with the sequence-file access. This usually indicates a software error.
NO UUCP SERVER	A TCP/IP call was attempted, but there is no server for UUCP.
BAD UID	The <i>uid</i> cannot be found in the <i>/etc/passwd</i> file. The filesystem is in trouble, or the <i>/etc/passwd</i> file is inconsistent.
BAD LOGIN_UID	The <i>uid</i> cannot be found in the <i>/etc/passwd</i> file. The filesystem is in trouble, or the <i>/etc/passwd</i> file is inconsistent.
ULIMIT TOO SMALL	The ulimit for the current user process is too small. File transfers may fail, so transfer will not be attempted.
BAD LINE	There is a bad line in the <i>Devices</i> file; there are not enough arguments on one or more lines.
FSTAT FAILED IN EWRDATA	There is something wrong with the Ethernet media.
SYSLST OVERFLOW	An internal table in <i>gename.c</i> overflowed. A big or strange request was attempted.
TOO MANY SAVED C FILES	An internal table in <i>gename.c</i> overflowed. A big or strange request was attempted.
RETURN FROM FIXLINE IOCTL	An <i>ioctl</i> , which should never fail, failed. There is likely a system driver problem.
PERMISSIONS file: BAD OPTION	There is a bad line or option in the <i>Permissions</i> file. Fix it immediately.
BAD SPEED	A bad line-speed appears in the <i>Devices</i> or <i>Systems</i> file (Class field).

Table 8-4 Assert Error Messages (**continued**)

Error Message	Description/Action
PKCGET READ	The remote station probably hung up. No action is required.
PKXSTART	The remote station aborted in a nonrecoverable way. This message can generally be ignored.
SYSTAT OPENFAIL	There is a problem with the modes of <i>/var/spool/uucp/.Status</i> , or there is a file with bad modes in the directory.
TOO MANY LOCKS	There is an internal problem.
XMV ERROR	There is a problem with some file or directory. The problem is likely caused by the spool directory, because the modes of the destinations should have been checked before this process was attempted.
CAN'T FORK	An attempt to fork and execute failed. The current job should not be lost, but will be attempted later (<i>uuxqt</i>). No action need be taken.

STATUS Error Messages

Status error messages are stored in the */var/spool/uucp/.Status* directory. This directory contains a separate file for each remote station that your station attempts to communicate with. These files contain status information on the attempted communication, indicating whether it was successful or not. Table 8-5 lists the most common error messages that can appear in these files.

Table 8-5 STATUS Error Messages

Error Message	Description/Action
OK	System status is normal
NO DEVICES AVAILABLE	There is currently no device available for the call. Make sure that there is a valid device in the <i>Devices</i> file for the particular station. Check the <i>Systems</i> file for the device to be used to call the station.
WRONG TIME TO CALL	A call was placed to the station at a time other than that specified in the <i>Systems</i> file.

Table 8-5 STATUS Error Messages (**continued**)

Error Message	Description/Action
TALKING	Self-explanatory.
LOGIN FAILED	The login for the given station failed. The problem could be a wrong login and password, wrong number, a very slow station, or failure in getting through the Dialer-Token-Pairs script.
CONVERSATION FAILED	The conversation failed after successful startup. This situation usually means that one side went down, the program aborted, or the line (link) was dropped.
DIAL FAILED	The remote station never answered. The problem could be a bad dialer or the wrong phone number.
BAD LOGIN/MACHINE COMBINATION	The station called you with a login or station name that does not agree with the <i>Permissions</i> file. This could be an attempt to breach system security.
DEVICE LOCKED	The calling device to be used is currently locked and in use by another process.
ASSERT ERROR	An ASSERT error occurred. Check the <i>/var/spool/uucp/.Admin/errors</i> file for the error message and refer to “Other UUCP Files” on page 232.
SYSTEM NOT IN Systems	The station is not in the <i>Systems</i> file.
CAN'T ACCESS DEVICE	Typically, this message means that the permissions on the device file (<i>/dev/tty*</i>) are not set correctly. Some programs set these permissions, and if terminated abnormally, do not reset them to correct states. Also, check the appropriate entries in the <i>Systems</i> and <i>Devices</i> files.
DEVICE FAILED	The attempt to open the device failed.
WRONG MACHINE NAME	The called station is reporting a name different from the one expected.
CALLBACK REQUIRED	The called station requires that it call your computer back to start a connection.

Table 8-5 STATUS Error Messages **(continued)**

Error Message	Description/Action
REMOTE HAS A LCK FILE FOR ME	The remote site has a LCK file for your computer. The remote station could be trying to call your computer. If they have an older version of UUCP, the process that was talking to your station may have failed earlier, leaving the LCK file. If the remote site has the new version of UUCP and they are not communicating with your computer, then the process that has a LCK file is hung.
REMOTE DOES NOT KNOW ME	The remote computer does not have the node name of your computer in its <i>Systems</i> file.
REMOTE REJECT AFTER LOGIN	The ID used by your computer to log in does not agree with what the remote computer was expecting.
REMOTE REJECT, UNKNOWN MESSAGE	The remote computer rejected the communication with your computer for an unknown reason. The remote computer may not be running a standard version of UUCP.
STARTUP FAILED	Login succeeded, but initial handshake failed.
CALLER SCRIPT FAILED	The problem indicated by this message is usually the same as that indicated by <code>DIAL FAILED</code> . However, if it occurs often, suspect the caller script in the <i>Dialers</i> file. Use <code>uutry</code> to check the caller script.

IRIX sendmail

This chapter describes IRIX sendmail, a facility for routing mail across an intranet. This chapter is for system administrators who set up and maintain the mail system on a station or network. It provides the information necessary for a straightforward implementation of sendmail. The following topics are covered:

- “About the Mail System” on page 254
- “Overview of sendmail” on page 255
- “Structure of sendmail” on page 257
- “About the sendmail Components” on page 259
- “Aliases Database for sendmail” on page 265
- “Configuring sendmail” on page 269
- “Managing sendmail” on page 275
- “About sendmail MX Records” on page 279

IRIX sendmail supports Internet Version 6 (IPv6). All IPv6 addresses must be prefixed by `IPv6:`. If you specify IPv6 addresses in your access database, local-host-names file, and/or relay-domains file, the files may need to be updated.

For more information on IPv6, see “Internet Protocol Version 6” on page 18.

For additional information on IRIX sendmail, see Appendix B, “IRIX sendmail Reference”.

About the Mail System

The mail system is a group of programs that you can use to send messages to and receive messages from other users on the network. You can send mail through either UUCP or TCP/IP. The IRIX operating system uses Netscape Mail, System V */bin/mail*, 4.3BSD */usr/sbin/Mail*, and the sendmail application for its mail implementation.

The process of delivering mail involves these elements:

User Interface The user interface creates new messages and reads, removes, and archives received messages. Netscape Mail, System V */bin/mail*, and 4.3BSD */usr/sbin/Mail* are the user interfaces provided with IRIX. Man pages are available to fully describe the features of these interfaces, and Netscape has an online help system.

Mail

Submission When a mail is sent locally from a Mail User Agent (MUA), the mail is handled by the Mail Submission Agent (MSA). The MSA is responsible for communicating the mail to the Mail Transfer Agent (MTA), that is, the local `sendmail` daemon running on the local host.

Note: This requires the implementation of special mail account. For information, see "sendmail Accounts".

Mail Routing A mail router examines each message and routes it through the network to the appropriate station. The sendmail program not only routes messages, but also formats them appropriately for their recipient stations.

Mail Transfer A mail transfer program transmits messages from one station to another. The sendmail application implements the Simple Mail Transfer Protocol (SMTP) over TCP/IP. For TCP/IP mail, sendmail acts as an integrated routing and transfer program. In all cases, mail transfer has a counterpart: mail reception. In most cases, a single program provides both functions. UUCP is a mail transfer program that uses its own protocols and runs over serial lines.

Mail Delivery A mail delivery program deposits mail into a data file for later perusal by a user or another program. The */bin/mail -d* program delivers local mail.

After you compose a message by using Netscape Mail, `/bin/mail`, or `/usr/sbin/Mail`, the message is sent to sendmail, which attempts to determine the destination of the message. The sendmail application either calls `/usr/bin/mail.local` (for mail to a user on the local station) or passes the message to the appropriate mail transfer program (for mail to a user on a remote station).

When sendmail receives a message from another station, it analyzes the recipient address; then, it either calls `/usr/bin/mail.local` to complete the delivery if the local station is acting as a relay, or passes the message to the mail transfer program. For TCP/IP SMTP, sendmail also performs the mail transfer.

When you send a mail message on a network that uses TCP/IP, several layers of network software are involved. Figure 9-1 shows the layers of TCP/IP mail network software.




Figure 9-1 Layers of TCP/IP Mail Software

Overview of sendmail

The Transmission Control Protocol (TCP) layer supports SMTP, which sendmail uses to transfer mail to other TCP/IP stations. The sendmail application is responsible for calling local delivery programs, mail routing, and TCP/IP mail transfer; it may also call other mail transfer programs. For example, sendmail uses the UUCP transmission program to handle messages sent to UUCP stations.

The sendmail implementation features aliasing, forwarding, automatic routing to network gateways, and flexible configuration.

In a simple network, each node has an address, and resources can be identified with a host-resource pair. For example, a mail system can refer to users with a host-user-name

pair. Station names and numbers must be administered by a central authority, but user names can be assigned locally to each station.

In an intranet, multiple networks with different characteristics and management must communicate. In particular, the syntax and semantics of resource identification change. You can handle certain simple cases by using improvised techniques, such as providing network names that appear local to stations on other networks. However, the general case is extremely complex. For example, some networks require point-to-point routing, which simplifies the database update problem, because only adjacent stations are entered into the system tables; others use end-to-end addressing. Some networks use a left-associative syntax; others use a right-associative syntax, causing ambiguity in mixed addresses.

Intranet standards seek to eliminate these problems. Initially, these standards proposed expanding the address pairs to address triples, consisting of *network*, *station*, *resource*. Network numbers must be universally agreed upon; stations can be assigned locally on each network. The user-level presentation was quickly expanded to address domains, composed of a local resource identification and a hierarchical domain specification with a common static root, as defined in RFC 1034. The domain technique separates the issue of physical versus logical addressing. For example, an address of the form `jane@iris1.company.com` describes only the logical organization of the address space.

The sendmail application bridges the gap between the world of totally isolated networks that know nothing of each other and the clean, tightly coupled world of unique network numbers. This application can accept old arbitrary address syntaxes, resolving ambiguities by using heuristics specified by the network administrator, as well as domain-based addressing. It helps guide the conversion of message formats between disparate networks. In short, sendmail is designed to assist a graceful transition to consistent intranet addressing schemes. This version of sendmail has the ability to connect to servers running on named sockets.

System Organization

The design goals for the sendmail application included the following:

1. Message delivery should be reliable, guaranteeing that every message is correctly delivered or at least brought to the attention of a human for correct disposal; no message should ever be completely lost.
2. Existing software should be used to do actual message delivery whenever possible.

3. The sendmail application should be easy to expand to fairly complex environments.
4. Configuration should not be compiled into the code.
5. The sendmail application should let various groups maintain their own mailing lists, and let individuals specify their own forwarding, without modifying the station's alias file.
6. Each user should be able to specify the mailer to execute to process mail being delivered. This feature allows users with specialized mailers that use a different format to build their environments without changing the system, and facilitates specialized functions (such as returning an "I am on vacation" message).
7. To minimize network traffic, addresses should be batched to a single station where possible, without assistance from the user.

Structure of sendmail

Figure 9-2 illustrates the sendmail system structure that is based on the original design goals for sendmail.




Figure 9-2 sendmail System Structure

The sendmail application neither interfaces with the user nor does actual mail delivery. Rather, it collects a message generated by a user agent program such as Berkeley *Mail*, edits the message as required by the destination network, and calls appropriate mailers to do mail delivery or queuing for network transmission. The exception is mail sent to a file; in this case, sendmail delivers the mail directly. This discipline allows the insertion of new mailers at minimum cost.

Because some of the senders may be network servers and some of the mailers may be network clients, sendmail can be used as an intranet mail gateway. The sendmail application supports multiple queues.

For security reasons, sendmail is no longer implemented as a `setuid-root` program. Rather, it is implemented as a `setgid-sgismmsp` program. The sendmail functionality is now split between MTA and MSA programs and requires new account and group categories both named `sgismmsp`. See “sendmail Accounts” on page 260.

About the sendmail Components

Understanding the sendmail programs requires understanding a variety of components. Some of these components are daemons, scripts, files, and commands. This section describes the various sendmail components:

- “sendmail Daemon” on page 259
- “sendmail Accounts” on page 260
- “sendmail Scripts” on page 260
- “sendmail Related Files and Directories” on page 262

sendmail Daemon

For sendmail to process incoming mail, several daemons must be running. The sendmail daemon is the sendmail program with specific flags. (Appendix B describes the *sendmail* command-line flags in detail.) The daemon is automatically started by the `/etc/init.d/mail` script at station startup. Controlling parameters are read from `/etc/config/sendmail.options` file. If this file does not exist, use the following command:

```
/usr/lib/sendmail -bd -q15m
```

The `-bd` flag causes *sendmail* to run in daemon mode. The `-q15m` flag causes sendmail to fork a subdaemon for queue processing every 15 minutes. The `-bd` and `-q` flags can be combined in one call. *sendmail* runs the queue(s) in the background when invoked with the `-q` option unless the `-qf` option or `-v` are used.

The `/etc/init.d/mail` script also starts the following daemon:

```
/usr/lib/sendmail -L sendmail.mspqueue -Ac -q15m
```

This daemon is responsible for clearing the mail queued by the Mail Submission Agent. It does not listen to port 25. It is invoked with the **-q15m** argument read from `/etc/config/sendmail.options` file.

sendmail Accounts

The `sendmail` program requires the following entries in the `/etc/passwd` and `/etc/group` files. When you upgrade to IRIX 6.5.19 (and later), these entries are added automatically to the `/etc/passwd` and `/etc/group`. If `/etc/shadow` password exists, it also updates this file using the `pwconv` command see the `pwconv(1M)` man page for more information.

If these entries need to be recreated, perform the following command as root:

```
/etc/mail/passwd.add
```

Add entries to the appropriate files as follows:

- Add the following entry to `/etc/passwd`:

```
sgismmsp:*LK*:25:25:Sendmail Mail Submission
Account:/dev/null:/dev/null
```
- Add the following entry to `/etc/group`:

```
sgismmsp:*:25:sgismmsp
```
- Add the following entry to `/etc/passwd.sgi`:

```
sgismmsp:noshow
```

sendmail Scripts

Several scripts provided with your system perform common functions in `sendmail`. Use the following scripts whenever possible, because they have been tested and are known to perform the task correctly:

- `"/etc/init.d/mail sendmail Script"` on page 261
- `"/usr/etc/configmail sendmail Script"` on page 261

`/etc/init.d/mail sendmail Script`

Under rare circumstances, a user may need to stop or start the sendmail daemon manually. For example, to implement changes to the configuration file, you must stop all running sendmail processes and restart the sendmail daemon before the new configuration will take effect. To simplify the task of starting and stopping sendmail, IRIX provides a shell script called `/etc/init.d/mail`.

This script takes a single argument, either **start** or **stop**, which starts or stops the sendmail daemon respectively. You must be superuser (root) to use this script. For example, to stop sendmail, use the following command:

```
/etc/init.d/mail stop
```

When `/etc/init.d/mail` is called with the **start** argument, it verifies the existence and permissions of various sendmail related files and directories (see “sendmail Related Files and Directories” on page 262). If a required component such as the `/var/spool/mqueue` directory is missing, the script creates it. For more complex components, such as `/etc/aliases`, the script exits with a message.

If the `sendmail_cf chkconfig` flag is on, new `/etc/sendmail.cf` and `/etc/mail/submit.cf` files are generated from the `/etc/sendmail.mc` and `/etc/mail/submit.mc` files, respectively.

When the `/etc/init.d/mail` script is called with the **stop** argument, it kills all running `sendmail` processes with a SIGTERM signal.

Note: Station start-up includes an automatic call to the `/etc/init.d/mail` script with the **start** argument. If station start-up runs in verbose mode (that is, `/etc/chkconfig` verbose on), the following message displays, verifying that sendmail has been started:

```
Mailer daemons: sendmail -bd -q15m ... sendmail.mspqueue -q15m
```

For more information, examine the `/etc/init.d/mail` script.

`/usr/etc/configmail sendmail Script`

The `/usr/etc/configmail` script provides an interface between command line input and the `sendmail.mc` file. The `configmail` parameters are maintained in the

sendmail.mc file. The *configmail* script migrates *configmail* parameter changes to the *sendmail.mc* file.

Historically, the *configmail* script has been used as the primary way to manage the *sendmail.cf* configuration file. You can use *configmail* to generate the */etc/sendmail.cf* file from the parameters that are stored in the *sendmail.mc* file by using the following command:

```
configmail mc2cf
```

Because the */etc/sendmail.mc* file offers more flexibility than *configmail*, you should modify *sendmail.mc* directly. Direct editing of the */etc/sendmail.cf* file is not supported.

A new */etc/mail/submit.cf* file is generated from the */etc/mail/submit.mc* file. Direct editing of the */etc/mail/submit.cf* file is not supported.

sendmail Related Files and Directories

The sendmail configuration files and directories are as follows:

- “*/etc/sendmail.cf* Configuration File” on page 263
- “*/etc/mail/submit.mc* and */etc/mail/submit.cf*” on page 263
- “*/etc/mail* sendmail configuration directory” on page 263
- “*/etc/sendmail.hf* Help File” on page 264
- “*/etc/sendmail.st* Statistics File” on page 264
- “*/etc/aliases* Aliases File” on page 264
- “*/var/spool/mqueue* Mail Queue Directory” on page 264
- “*/var/spool/clientmqueue*” on page 265
- “*/usr/lib/sendmail.cf_m4*” on page 265
- “*/var/mail* Incoming Mail Directory” on page 265

The sendmail application also can be tailored with flags and options from the command line. For more information, see the *sendmail(8)* man page.

/etc/sendmail.cf Configuration File

At the heart of the sendmail program is the sendmail configuration file, */etc/sendmail.cf*. The *sendmail.cf* file is an ASCII file that contains most of the configuration information and is read at run time. This file encodes options, header declarations, mailer declarations, trusted user declarations, message precedences, address-rewriting rules, macro definitions, and class definitions.

To successfully set up sendmail, the mail administrator must know which *sendmail.cf* macros and variables to change.

/etc/mail/submit.mc and /etc/mail/submit.cf

The *submit.mc* file is the configuration file used by the MSA portion of sendmail. Normally, you are not required to change it. It is written for the m4 processor and is also handled by the `configmail mc2cf` command. The result of this conversion is the *submit.cf* file.

/etc/sendmail.mc Configuration File

The *sendmail.mc* file contains a series of commands that will be used to build a *sendmail.cf* file. The *sendmail.mc* file is also used to enable various features of the *sendmail.cf* file. The *sendmail.mc* file is written for the m4 macro processor. Thus, all quoted strings use both open and close quotes ('example'). For details on configuring *sendmail.mc*, see "Configuring sendmail" on page 269.

To generate a *sendmail.cf* file from a *sendmail.mc* file, use the following command:

```
configmail mc2cf
```

/etc/mail sendmail configuration directory

The sendmail program expects **all** configuration information to reside in the */etc/mail* directory, making it easier to access all sendmail configuration-related files. Because of this, */etc/init.d/mail* script will make the appropriate symbolic links from old locations onto the */etc/mail* location. For example, */etc/mail/sendmail.cf* is linked to */etc/sendmail.cf*. The details of these operations can be found in the mail startup script in the function `make_symlinks_893`. Also note that no files in */etc/mail* directory are overwritten.

/etc/sendmail.hf Help File

The *sendmail.hf* file is the Simple Mail Transfer Protocol (SMTP) help file. It contains some brief information about the various SMTP commands.

/etc/sendmail.st Statistics File

The *sendmail.st* file collects statistics related to sendmail. By default, the file is not present. You can create the file using the *touch* command. If the file is present, sendmail automatically updates the file with relevant sendmail statistics.

/etc/aliases Aliases File

The *aliases* file contains the text form of the alias database used by the sendmail program. The alias database contains aliases for local mail recipients. For example, the following alias delivers mail addressed to jd on the local station to johndoe@company.com:

```
jd: johndoe@company.com
```

Note: The *newaliases* program must be run after modifying the alias database file. For more information about building the alias database, see “Building the sendmail Aliases Database” on page 266.

/var/spool/mqueue Mail Queue Directory

The mail queue, */var/spool/mqueue*, is the directory in which the mail queue and temporary files reside. The messages are stored in various queue files that exist under the */var/spool/mqueue* directory. Queue files take the following forms:

- *qf**—control (queue) files for messages
- *df**—data files
- *tf**—temporary files
- *nf**—a file used when a unique ID is created
- *xf**—transcript file of the current session

Normally, a sendmail subdaemon processes the messages in this queue periodically, attempting to deliver each message. (The */etc/init.d/mail* script starts the sendmail

daemon so that it forks a subdaemon every 15 minutes to process the mail queue.) Each time sendmail processes the queue, it reads and sorts the queue, then attempts to run all jobs in order.

Note: Any mail queue files generated this way cannot be read by previous versions of sendmail.

`/var/spool/clientmqueue`

This is the mail queue directory used by the MSA portion of sendmail. Files in this directory have the same format as described for the `/var/spool/mqueue` directory. By default, as defined in the `configmail.options` file, this directory is cleaned every 15 minute by a second daemon invocation of sendmail (for more information, see “sendmail Daemon” on page 259).

`/usr/lib/sendmail.cf_m4`

The `/usr/lib/sendmail.cf_m4` directory contains the data that is used to generate the `sendmail.cf` file. The contents of this directory are similar to the `cf` subdirectory from the sendmail source distribution. The `sendmail.mc` file selects settings and features implemented by the files in the `/usr/lib/sendmail.cf_m4` directory. Use the following command to generate a new `sendmail.cf` file from these files:

```
configmail mc2cf
```

`/var/mail Incoming Mail Directory`

The `/var/mail` directory houses all incoming mail. Each user on a local station receives his or her mail in a file in the `/var/mail` directory. For example, the user `guest` receives mail in the `/var/mail/guest` file.

Aliases Database for sendmail

This section explains how the aliases database for sendmail runs, and indicates possible problems and errors. Sections are as follows:

- “Building the sendmail Aliases Database” on page 266

- “Testing the sendmail Aliases Database” on page 268
- “sendmail Alias Database Problems” on page 268
- “sendmail List Owners” on page 269

The aliases database is a NEWDB database that contains mail aliases that are used by the sendmail program. The text form of the database is maintained in the */etc/aliases* file. The aliases are of the following form:

```
name: name1 [ , name2, ... ]
```

For example, the following command delivers mail addressed to jd to johndoe@company.com:

```
jd: johndoe@company.com
```

Note: Only the local part of an address can be aliased. For example, the following command is wrong and does not have the desired effect:

```
jd@big.university.edu:jd@company.com
```

The sendmail application consults the alias database only after deciding that the message (as originally addressed) should be delivered locally, and after it has rewritten the address to contain only the local part.

An alias continuation line must start with a space or a tab. Blank lines and lines beginning with the number sign (#) are treated as comments.

Building the sendmail Aliases Database

At startup, sendmail automatically uses the *newdb* library to process the */etc/aliases* file into the */etc/aliases.db* files. Using these files to resolve aliases improves performance.

When using the NEWDB version of the database to maintain aliases, sendmail uses the */etc/alias.db* file. When building the NEWDB version of the database, sendmail checks the left-hand side of each entry to make sure that it is a local address. Then sendmail issues a warning for each entry in */etc/aliases* with a non-local left-hand side. Such entries are not entered into the NEWDB version of the database.

Previous versions of IRIX sendmail looked for the keyword `++` in the `/etc/alias` file and enabled NIS lookups for you. You must now do this explicitly. By default, IRIX sendmail uses the `/etc/aliases` database and the NSD map, `mail.aliases`, for lookups. If you have configured nsd to use NIS, you do not need to do anything more. If you have not done this configuration, you can place the following line in the `sendmail.mc` file:

```
define (ALIAS_FILE',/etc/aliases,nsd:mail.aliases') dnl
```

If the configuration or the command line specifies the **D** option, sendmail automatically tries to rebuild the alias database when it is out of date.

The sendmail application rebuilds the alias database if either of the following conditions exists:

- The NEWDB version of the database is in mode 666.
- sendmail is running `setuid` to root.

Note: Auto-rebuild can be dangerous on heavily loaded stations with large alias files. If it takes more than 5 minutes to rebuild the database, there is a chance that several processes will start the rebuild process simultaneously.

The `newaliases` program is used to rebuild the NEWDB version of the `aliases` database. This program must be run any time the text version of the `aliases` file is modified. If `newaliases` is not run after making changes to the `aliases` file, the changes are not incorporated into the NEWDB `alias` database and are not seen by the sendmail program.

To rebuild the NEWDB version of the database without restarting sendmail, enter the following command:

```
newaliases
```

Executing the *newaliases* command is equivalent to giving *sendmail* the **-bi** flag, as in the following example:

```
/usr/lib/sendmail -bi
```

Testing the sendmail Aliases Database

You can test the *aliases* database with the **-bv** flag of the *sendmail* program. For more details, see Appendix B.

sendmail Alias Database Problems

Problems can occur with the *aliases* database, especially if a sendmail process accesses the NEWDB version before it is completely rebuilt. Two circumstances can cause this problem:

- One process accesses the database while another process is rebuilding it.
- The process rebuilding the database dies because it has been killed, or a station crash has occurred before completing the rebuild.

The sendmail application has two techniques for trying to relieve these problems. First, to avoid the problem of a partially rebuilt database, sendmail ignores interrupts while rebuilding the database. Second, at the end of the rebuild it adds an alias of the following form (which is not normally legal):

```
@: @
```

Before sendmail accesses the database, it ensures that this entry exists. For this action to occur, the configuration file must contain the **-a** option.

If the @:@ entry does not exist, sendmail waits for it to appear. After the specified waiting period elapses, sendmail itself forces a rebuild. For this action to occur, the configuration file must include the **D** option. If the **D** option is not specified, a warning message is generated and sendmail continues.

sendmail List Owners

If an error occurs when mail is sent to a certain address (*x*, for example), sendmail looks for an alias of the following form to receive the errors:

```
owner-x
```

This scheme is typically useful for a mailing list where a user mailing to the list has no control over the maintenance of the list itself; in this case the list maintainer would be the owner of the list. For example, the following would cause `jd@1company.com` to get the error that occurs when someone sends mail to `unix-hackers`, and sendmail finds the phony user `nosuchuser` on the list.

```
unix-hackers: jd@company1.com, ed@big.university.edu,nosuchuser,  
jane@company2.com  
owner-unix-hackers: jd@company1.com
```

Configuring sendmail

This section provides sendmail configuration information. It does not attempt to be a complete reference for programming the `sendmail.cf` file and all of its features. For details of sendmail configuration, see the following resources:

- *sendmail, 2nd Edition* by Bryan Costales & Eric Allman. This book, considered the definitive sendmail authority, is available from O'Reilly & Associates Inc. (ISBN: 1-56592-222-0) (URL: <http://www.oreilly.com/catalog/sendmail2>)
- The *sendmail.org* website (<http://www.sendmail.org>) has a large amount of sendmail information.
- The `/usr/lib/sendmail.cf_m4/README` file includes detailed information on the `sendmail.mc` file and the m4 macro processor features and configuration. This file includes define statements that you can make to change parameters.

sendmail.cf File

The `sendmail.cf` file is read in real time by each copy of sendmail. There is no frozen version of this file, which previous versions of IRIX sendmail supported. This section does not provide details of the `sendmail.cf` file. Instead, it describes how to configure and generate it by using the `sendmail.mc` file

sendmail.mc and sendmail.cf Files

The *sendmail.mc* file is a set of definitions, configurations, and commands to enable various features. Use this file to create the *sendmail.cf* file from the prototype files located in the */usr/lib/sendmail.cf_m4* directory.

To create a *sendmail.cf* file from the definitions in *sendmail.mc* file, use the following command:

```
configmail mc2cf
```

By using the defines and features in the following sections, you can implement a number of complex features, including, but not limited to those in the following list:

- Masquerading and relaying
- Controlling anti-spam configuration
- Adding new mailers or rulesets
- Using non-smtp based configurations
- Using mailer tables
- Using userdb to map full names

The */usr/lib/sendmail.cf_m4/README* file describes the sendmail configuration files and explains how to create a *sendmail.cf* file for use with sendmail. This section shows the headings and describes the contents of the README file and covers the following topics:

- “m4 Macro Processor” on page 271
- “File Locations” on page 272
- “Operating System Type” on page 272
- “Domains” on page 272
- “Mailers” on page 272
- “Special Features” on page 272
- “Using UUCP Mailers” on page 272
- “Tweaking Rulesets” on page 273
- “Masquerading and Relaying” on page 273

- “Anti-SPAM Configuration Control” on page 273
- “Adding New Mailers or Rulesets” on page 273
- “Adding New Mail Filters” on page 273
- “Queue Group Definitions” on page 273
- “Accepting Mail for Multiple Names” on page 274
- “Using Mailertables” on page 274
- “Using the User Database To Map Full Names” on page 274
- “Miscellaneous Special Features” on page 274
- “Security Notes” on page 274
- “Tweaking Configuration Options” on page 275
- “Message Submission Program” on page 275
- “Format of Files and Maps” on page 275
- “Directory Layout” on page 275
- “Administrative Details” on page 275

m4 Macro Processor

The sendmail application uses the m4 macro processor (simply known as m4) to “compile” the configuration files. M4 is stream-based, that is, it does not understand about lines. For this reason, in some places you may see the word “dnl”, which stands for “delete through newline”; essentially, it deletes all characters starting at the “dnl” up to and including the next newline character. In most cases, sendmail uses this only to avoid unnecessary blank lines in the output. Other important directives are `define(A, B)` which defines the macro “A” to have value “B”. Macros are expanded as they are read, so one normally quotes both values to prevent expansion. For example, `define(SMART_HOST, 'smart.foo.com')` One word of warning: M4 macros are expanded even in lines that appear to be comments. For example, if you have `See FEATURE('foo')` above it will not do what you expect, because the `FEATURE('foo')` will be expanded. This also applies to `And` then define the `$X` macro to be the return address because “define” is an M4 keyword. If you want to use them, surround them with directed quotes, ‘like this’. For more information on M4, see `/usr/lib/sendmail.cf_m4/README` file.

File Locations

This section of the README file describes the `/etc/mail` configuration directory for sendmail related files. For a list of files, see the `/usr/lib/sendmail.cf_m4/README` file.

Operating System Type

You must define an operating system environment, or the configuration file build will fail. There are several environments available; look at the `ostype` directory for the current list. This macro changes things like the location of the alias file and queue directory. For more information on the `ostype` directory, see the `/usr/lib/sendmail.cf_m4/README` file.

Domains

You will probably want to collect domain-dependent defines into one file, referenced by the `DOMAIN` macro. For example, the Berkeley domain file includes definitions for several internal distinguished hosts. For more information on the `DOMAIN` macro, see the `/usr/lib/sendmail.cf_m4/README` file.

Mailers

This section of the README file describes the mailers available with this release. As a general rule, put the `MAILER` definitions last in your `.mc` file. For more information on sendmail mailers, see the `/usr/lib/sendmail.cf_m4/README` file.

Special Features

Special features can be requested using the `FEATURE` macro. For more information on the `FEATURE` macro, see the `/usr/lib/sendmail.cf_m4/README` file.

Using UUCP Mailers

There are four UUCP mailers available. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Tweaking Rulesets

For more complex configurations, you can define special rules. The macro `LOCAL_RULE_3` introduces rules that are used in canonicalizing the names. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Masquerading and Relaying

You can have your host masquerade as another. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Anti-SPAM Configuration Control

The primary anti-spam features available in sendmail are:

- Relaying is denied by default
- Better checking on sender information
- Access database
- Header checks

For more information on anti-SPAM features, see the `/usr/lib/sendmail.cf_m4/README` file.

Adding New Mailers or Rulesets

Sometimes you may need to add entirely new mailers or rulesets. They should be introduced with the constructs `MAILER_DEFINITIONS` and `LOCAL_RULESETS`, respectively. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Adding New Mail Filters

Sendmail supports mail filters to filter incoming SMTP messages according to the *Sendmail Mail Filter API* documentation. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Queue Group Definitions

In addition to the queue directory (which is the default queue group called `mqueue`), sendmail can deal with multiple queue groups, which are collections of queue directories

with the same behavior. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Accepting Mail for Multiple Names

If your host is known by several different names, you need to augment class {w}. This is a list of names by which your host is known, and anything sent to an address using a host name in this list will be treated as local mail. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Using Mailertables

To use mailertables feature, you have to create an external database containing the routing information for various domains. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Using the User Database To Map Full Names

The user database was not originally intended for mapping full names (for example., Eric.Allman => eric), but some sites are using it that way. It is recommended that you set up aliases for this purpose instead since you can specify multiple alias files. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Miscellaneous Special Features

Sometimes it is convenient to merge configuration on a centralized mail machine, for example, to forward all root mail to a mail server. In this case it might be useful to be able to treat the root addresses as a class of addresses with subtle differences. You can do this using plussed users. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Security Notes

For a list of security considerations, see the `/usr/lib/sendmail.cf_m4/README` file.

Tweaking Configuration Options

There are a large number of configuration options that you do not normally need to change. For a list of M4 variables that can tune your configuration, see the `/usr/lib/sendmail.cf_m4/README` file.

Message Submission Program

The purpose of the message submission program (MSP) is explained in the sendmail security section. For a list of caveats and a few hints how tune the default configuration (which is installed as `submit.cf`), see the `/usr/lib/sendmail.cf_m4/README` file.

Format of Files and Maps

Files that define classes, for example, `F{classname}`, consist of lines each of which contains a single element of the class. For an example of `/etc/mail/local-host-names` file, see the `/usr/lib/sendmail.cf_m4/README` file.

Directory Layout

This section describes the subdirectories in the directory. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Administrative Details

This section of the README file describes certain internal parts of the `sendmail.cf` file. If you are trying to modify your sendmail configuration, read this section carefully. For more information, see the `/usr/lib/sendmail.cf_m4/README` file.

Managing sendmail

These aspects of managing the sendmail environment are covered in the following sections:

- “Starting the sendmail Daemon” on page 276
- “Listing the sendmail Message Queue” on page 276

- “Forcing the sendmail Message Queue” on page 276
- “Redirecting Mail With the .forward File” on page 278

Further aspects of managing the sendmail environment are covered in Appendix B, “IRIX sendmail Reference”.

Starting the sendmail Daemon

After customizing the *sendmail.cf* files and modifying the *aliases* database, you are ready to start sendmail.

By default, IRIX automatically starts sendmail at station startup by using the shell script */etc/init.d/mail*. However, if you are configuring and testing sendmail and do not want to reboot the station, you can run the */etc/init.d/mail* script manually. You should always use the *mail* script to stop and start sendmail. It processes and checks sendmail related files and programs in the correct order.

Start the sendmail daemon:

```
/etc/init.d/mail start
```

If you need to stop sendmail, enter the following command:

```
/etc/init.d/mail stop
```

Listing the sendmail Message Queue

You can list the contents of the queue by using the *mailq* command or by specifying the **-bp** flag on the *sendmail* command. The list includes a listing of the queue IDs, the size of each message, the date the message entered the queue, and the sender and recipients.

Forcing the sendmail Message Queue

The **-q** flag (with no value) forces sendmail to process the queue. It is sometimes useful to use the **-v** flag (verbose) also when running the queue manually, as follows:

```
/usr/lib/sendmail -q -v
```

In verbose mode, sendmail displays the SMTP chatter with other stations as well as messages indicating any delivery errors and final message disposition.

Because of the locking algorithm, it is impossible for one job to freeze the queue. However, an uncooperative recipient station or a program recipient that never returns can consume many station resources. Unfortunately, there is no way to resolve this situation without violating the SMTP protocol that sendmail uses.

In some cases, if a major station goes down for a couple of days, a prohibitively large queue may be created. As a result, sendmail spends an inordinate amount of time sorting the queue. You can remedy this situation by moving the queue to a temporary location and creating a new queue. The old queue can be run later when the offending station returns to service.

Use the following commands to move the entire queue directory. The mail queue should be owned by *root* and belong to the *mail* group.

1. Change to the `/var/spool` directory.
`cd /var/spool`
2. Give the old mail queue another filename.
`mv mqueue omqueue`
3. Make another directory for the new mail queue.
`mkdir mqueue`
4. Change permissions so that the directory is owned by root.
`chmod 755 mqueue`

Then kill the existing sendmail daemon (because it will still be processing in the old queue directory) and create a new daemon:

1. Stop the sendmail daemon that is currently running.
`/etc/init.d/mail stop`
2. Start a new sendmail daemon.
`/etc/init.d/mail start`
3. To run the old mail queue, use the following command:
`/usr/lib/sendmail -oQ/var/spool/omqueue -q`

The `-oQ` flag specifies an alternate queue directory, and the `-q` flag causes sendmail to run every job in the queue once and then return. Use the `-v` (verbose) flag to watch what is going on. It might be necessary to run the old mail queue a number of times before all of the messages can be delivered.

4. When the queue is finally emptied, the directory can be removed with the following command:

```
rmdir /var/spool/omqueue
```

Redirecting Mail With the `.forward` File

As an alternative to the alias database, users can put a file with the name `.forward` in their home directories. If the `.forward` file exists in a user's home directory, sendmail redirects mail for that user to the list of recipients in the file. The recipients are separated by commas or new lines. For example, if the home directory for user jane has a `.forward` file with the following contents, any mail arriving for jane is redirected to the specified accounts:

```
zippy@state.edu  
bongo@widgets.com
```

The `.forward` file also allows the user to redirect mail to files or programs. A `.forward` file with the following contents redirects any incoming messages to `jd@company.com`, appends a copy of the message to the file `/var/tmp/mail.log`, and pipes a copy of the message to `stdin` of the `/usr/bin/mymailer` program:

```
jd@company.com  
/var/tmp/mail.log  
| /usr/bin/mymailer
```

In general, file-type recipients must be writable by everyone. However, if sendmail is running as `root` and the file has `setuid` or `setgid` bits set, the message is written to the file.

Users can redirect mail to themselves in addition to sending it to other recipients. This feature is particularly useful if the users want to continue to receive mail in their own mailboxes while passing copies of each incoming message to some alternative destination. For example, if the home directory for user john contains a `.forward` file with the following contents:

```
\john, |/usr/sbin/vacation
```

The sendmail program behaves as follows:

- It sends each incoming message to john's regular mailbox; the backslash (\) preceding the name indicates that no further aliasing is to occur.
- It pipes a copy of each message to *stdin* of the */usr/sbin/vacation* program. (The vertical bar [|] is the standard UNIX pipe symbol.)

About sendmail MX Records

MX records are resource records in the BIND database. Each record contains the name of a target station, a preference level, and the name of an exchanger station that handles mail for the target station. (The exchanger station may be the target station itself.)

The BIND database can contain several MX records for each target station; the record with the lowest preference level is tried first.

MX records provide a way to direct mail to alternative stations. Using MX records lets you eliminate static routes from your sendmail configuration file.

For each destination station contacted by means of an IPC-type mailer (P=[IPC] in the mailer definition line), sendmail queries the DNS database for MX records associated with the destination station. If the MX query succeeds, mail will be routed through the appropriate exchanger station found among the returned MX records as described in RFC 974 and required by RFC 1123.

The result is that this version of sendmail and previous versions may use different methods to route mail to stations for which MX records are available. See Appendix B, "IRIX sendmail Reference", for information regarding mailer definitions.

SGI Infinite Network Bandwidth

This chapter describes SGI Infinite Network Bandwidth and covers the following topics:

- “SGI Infinite Network Bandwidth Overview” on page 282
Defines SGI Infinite Network Bandwidth and describes some advantages of using it.
- “Configuring a Network Stripe” on page 286
Provides instructions for configuring SGI Infinite Network Bandwidth on your system.
- “Verifying a Network Stripe” on page 289
Describes how to make sure that a remote stripe can be accessed.
- “Stripe Interface MTUs” on page 290
Describes how to set the Maximum Transmission Unit (MTU) size so Gigabit Ethernet and striping interfaces use jumbo frames.
- “Network Stripe Monitoring” on page 293
Describes tools you can use to monitor stripe performance.
- “Tuning a Network Stripe Interface” on page 294
Describes how to improve the performance of a network strip by tuning stripe parameters.
- “Application I/O Sizes” on page 299
Describes how to use SGI Infinite Network Bandwidth to achieve optimal I/O patterns for optimal application throughput.
- “Troubleshooting SGI Infinite Network Bandwidth Software” on page 300
Describes how to troubleshoot your SGI Infinite Network Bandwidth software.

SGI Infinite Network Bandwidth Overview

This section provides an overview of SGI Infinite Network Bandwidth and covers the following topics:

- “New Terminology Used in This Guide” on page 282
- “Definition of SGI Infinite Network Bandwidth” on page 283
- “Hardware Requirements” on page 283

New Terminology Used in This Guide

This section describes terminology associated with SGI Infinite Network Bandwidth.

- *Physical interface*
A network interface that represents an actual physical connection to an external network.
- *Logical or virtual interface*
A network interface that has no physical connections to an external network, but still provides a connection across a network made up of physical interfaces.
- *Tunnelling*
A method of encapsulating data of an arbitrary type inside another protocol header to provide a method of transport for that data across a network of a different type. A tunnel requires two endpoints that understand both the encapsulating protocol and the encapsulated data payload.
- *Link pair*
A link pair is defined as a pair of interconnected physical interfaces on different machines that can be used by a Network Stripe Interface. They are defined by the primary IPv4 addresses of the physical interfaces.
- *Network stripe interface*
A logical point to point network interface that uses tunnelling across multiple **link pairs** to send data between hosts. The data that is sent is distributed evenly ("striped") across all interfaces, therefore, allowing the logical interface to use the combined bandwidth of all the physical interfaces belonging to it. A **network stripe interface** that is made up of **n link pairs** is often referred to as an **n-stripe** or an **n-link stripe**.

Definition of SGI Infinite Network Bandwidth

SGI Infinite Network Bandwidth is very similar in concept to disk striping. Like disk striping, it is not a new concept. However, unlike disk striping, SGI Infinite Network Bandwidth network stripes are difficult to scale efficiently.

Striped disks are commonly used when the bandwidth that a single disk can provide to a filesystem is not enough for the required purpose. Multiple disks are striped together so that the filesystem blocks are spread evenly across all the disks in the stripe set so that data is retrieved from or stored to many disks in parallel. This greatly increases the bandwidth available to the filesystem.

SGI Infinite Network Bandwidth network stripes aggregate a number of physical network interfaces into a single virtual interface that is able to spread data evenly across all the physical interfaces in the stripe set. This allows multiple packets to be sent and received in parallel. This greatly increases the available bandwidth to the virtual interface.

Like a filesystem that is made up of a striped disk array, applications are able to make use of the increased bandwidth supplied by the network stripe interface without modifications. That is, the virtual interface is completely transparent to applications but provides significantly higher throughput.

In ideal conditions, a network stripe interface configured from a set of n identical interfaces will provide n -times the bandwidth of slowest interface that is part of the stripe set. In real world conditions, the number of packets that can be sent over a network stripe interface is generally limited by CPU speed.

Hardware Requirements

There are no particular hardware requirements for a network stripe. However, it is worth keeping these things in mind when deciding to use a network stripe, as follows:

- SGI Infinite Network Bandwidth has been designed for use with Ethernet networks. Therefore, using non-Ethernet devices may give unpredictable results.
- The network stripe driver is optimized for high bandwidth, low delay networks that have low packet loss and minimal packet reordering issues. Use on high delay networks and/or networks that suffer from excessive packet loss are is not recommended.

- The network stripe driver is optimized for homogeneous physical link interfaces. It works with different interface types, but it is limited in packet rate by the slowest interface, and limited in packet size by the interface with the smallest Maximum Transmission Unit (MTU). For example, striping a 100Mbit Ethernet and a Gigabit Ethernet interface will only result in a network stripe with the performance of 2x100Mbit Ethernet.
- Not all applications can benefit from a network stripe. If the application is latency bound, then a network stripe will not improve performance. If the application is not capable of using all the bandwidth available to a single interface on your machine, then a network stripe will not improve performance.

If you do have an application that is constrained by the bandwidth available to a single interface, for example, NFS, FTP, or rsync, then a network stripe may be a way of improving throughput.

- While SGI Infinite Network Bandwidth runs on any machine that is supported by IRIX 6.5.23, or later, the number and speed of the CPUs in the machines being used for stripe interfaces greatly impacts performance. Faster CPUs result in significantly better stripe performance.

For best results using Gigabit Ethernet network stripes, SGI recommends the following:

- Using an SGI Origin 300 or SGI Origin 3000 series system with a minimum of four 500 MHz processors.
- Using jumbo frames (9k MTU) on physical links where available for better performance and lower CPU utilization.

To configure the MTU size, see “MTU Size Settings” in the *SGI IRIS Release 2 Gigabit Ethernet Board User’s Guide*.

- For jumbo frames, only Gigabit Ethernet PCI-X network cards are supported. They appear as `tg` interfaces on your system. Striping with jumbo frames on `eg` hardware interface types is **not reliable** due to hardware limitations and is therefore **not recommended** or supported.

You can use the `hinv(1M)`, `ifconfig(1M)`, or `netstat(1)` commands to determine if your system has `tg` interfaces, as follows:

Using the `hinv(1M)` command:

```
root@mig127# hinv
2 500 MHZ IP35 Processors
CPU: MIPS R14000 Processor Chip Revision: 2.4
FPU: MIPS R14010 Floating Point Chip Revision: 2.4
Main memory size: 2048 Mbytes
```



```

Instruction cache size: 32 Kbytes
Data cache size: 32 Kbytes
Secondary unified instruction/data cache size: 2 Mbytes
Integral SCSI controller 0: Version QL12160, low voltage
differential
  Disk drive: unit 1 on SCSI controller 0
Integral SCSI controller 1: Version QL12160, low voltage
differential
IOC3/IOC4 serial port: tty3
IOC3/IOC4 serial port: tty4
Gigabit Ethernet: tg1, module 001c01, PCI bus 2 slot 1
Gigabit Ethernet: tg2, module 001c01, PCI bus 2 slot 2
Integral Fast Ethernet: ef0, version 1, module 001c01, pci 4
IOC3/IOC4 external interrupts: 1
USB controller: type OHCI

```

Using the `ifconfig -a(1M)` command:

```

root@mig127# ifconfig -a
ef0:
flags=8415c43<UP,BROADCAST,RUNNING,FILTMULTI,MULTICAST,CKSUM,DRVRLOC
K,LINK0,IPALIAS,IPV6>
  inet 128.162.246.127 netmask 0xffffffff00 broadcast
128.162.246.255
tg1: flags=8c85802<BROADCAST,MULTICAST,CKSUM,DRVRLock,LINKDOWN,IPALIA
S,HIGHBW,IPV6>
tg2: flags=8c04802<BROADCAST,MULTICAST,DRVRLock,IPALIAS,HIGHBW,IPV6>
cl0: flags=404800<MULTICAST,DRVRLock,IPALIAS>
lo0: flags=8001849<UP,LOOPBACK,RUNNING,MULTICAST,CKSUM,IPV6>
  inet 127.0.0.1 netmask 0xff000000

```

Using the `netstat -iq(1)` command:

```

root@mig127# netstat -iq

```

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	q	max	drop
ef0	1500	eagan-24	mig127.americas.	699398	0	561505	0	0	0	0
tg1*	1500	none	none	0	0	0	0	0	2048	0
tg2*	1500	none	none	0	0	0	0	0	0	0
cl0*	57344	none	none	0	0	0	0	0	64	0
lo0	32992	loopback	localhost	36871	0	36871	0	0	50	0

- Connecting the two sets of tg cards back-to-back. It is also possible to connect them through a switch that can support the bandwidth of multiple Gigabit Ethernet links simultaneously. If jumbo frames are being used, the switch must be able to support them as well.

Configuring a Network Stripe

To configure a network stripe on your system you need to use the `ifconfig(1M)` command. There are several new `ifconfig` parameters used to manipulate network stripe interfaces, as follows:

- `stripe laddr raddr`

This parameter adds a physical link to a stripe, where *laddr* is the primary address of the local physical interface the stripe will use, and *raddr* is the corresponding address of the remote physical interface to which the stripe will send packets and from which it will receive packets. If the stripe interface that is the target of this command does not exist, it is created as a side effect of running this command. If the stripe interface does already exist, it must be down for this command to succeed.

To add a link-pair to an `is` interface that is already up, the interface has to be brought down first as shown in the following example. In this example, two additional link-pairs are added to the stripe interface `is0`:

```
> ifconfig is0 down
> ifconfig is0 stripe 10.1.3.7 10.1.3.8
> ifconfig is0 stripe 10.1.2.7 10.1.2.8
```

- `stripe laddr`

This parameter removes a physical link pair from the current stripe set. The link pair that is removed is the one that matches the local address *laddr*. The stripe interface must be down for this command to succeed. If the link being removed is the last physical link on the stripe interface, the stripe interface will also be deleted as a side effect of this command.

In the following example, assume that `is0` already has 2-link pairs in its link set. To remove one (or both) of the link pairs from the stripe, first the `is0` interface has to be brought down; then `-stripe` option is used to remove the link-pairs from the stripe, as follows:

```
> ifconfig is0 down
> ifconfig is0 -stripe 10.1.2.7
> ifconfig is0 -stripe 10.1.3.7
```

- `stripelist`

The `stripelist` option lists the physical link pairs configured on a stripe interface. This information is also displayed when the `ifconfig` command is executed with the `-a` option and it encounters a stripe interface. This option only shows the link pairs that were present at the time the stripe interface was last brought up.

In the following example, the `stripelist` option lists all the link-pairs that belong to a striping interface `is0`.

```
> ifconfig is0 stripelist
           link local 10.1.2.7 remote 10.1.2.8
           link local 10.1.3.7 remote 10.1.3.8
```

Because a network stripe interface is a point-to-point interface, the `ifconfig` command requires you to give it both a source and destination address when you configure it.

A network stripe interface is created automatically when one of the following `ifconfig` configuration commands is executed. However, the stripe will not function until the stripe has both a source and destination address defined and it has at least one physical link configured.

The `ifconfig` parameter commands create a stripe interface either explicitly or implicitly, as follows:

- Explicitly by defining the source and destination addresses for the stripe interface, as follows:

```
> ifconfig is0 10.0.0.194 dst 10.0.0.195
```

- Implicitly by the adding the first physical link to a previously undefined stripe interface, as follows:

```
> ifconfig is0 stripe 192.168.1.194 192.168.1.195
```

The stripe will **not** be correctly configured until both of the above commands have been executed. The stripe must have both a source and destination address defined and it must have at least one physical link configured.

Note that the stripe interface will come up as soon as it has both a source and destination address defined. When the stripe interface is up, you cannot change the physical link pair configuration. Hence, if you define the interface first by giving it a source and destination address, you must mark it down before adding the physical link pairs, and then mark it up before it becomes active.

Procedure 10-1 Configuring a Simple Stripe Interface

To create a simple 1-stripe from scratch, perform these commands, as follows:

```
> ifconfig is0 10.0.0.194 dst 10.0.0.195
> ifconfig is0 down
> ifconfig is0 stripe 192.168.1.194 192.168.1.195
> ifconfig is0 up
```

The following command sequence is equivalent to the above sequence, but is less verbose:

```
> ifconfig is0 stripe 192.168.1.194 192.168.1.195
> ifconfig is0 10.0.0.194 dst 10.0.0.195
```

Procedure 10-2 Creating a 2-link stripe

The following steps describes how to create a 2-link stripe, as follows:

Machine A has two `tg` interfaces, `tg1` and `tg2`, with IP addresses 192.168.1.194 and 192.168.2.194, respectively. Machine B has two `tg` interfaces, `tg1` and `tg2`, with IP addresses 192.168.1.195 and 192.168.2.195, respectively.

To set-up a stripe with `tg1` and `tg2` interfaces, execute the following commands on machine A and B, as follows:

- On machine A create the stripe interface by defining its physical links, as follows:
Machine A > `ifconfig is0 stripe 192.168.1.194 192.168.1.195`
Machine A > `ifconfig is0 stripe 192.168.2.194 192.168.2.195`
- On machine B create the stripe interface by defining its physical links, as follows:
Machine B > `ifconfig is0 stripe 192.168.1.195 192.168.1.194`
Machine B > `ifconfig is0 stripe 192.168.2.195 192.168.2.194`
- On machine A define the stripe interface endpoint addresses and bring up the stripe, as follows:
Machine A > `ifconfig is0 10.0.0.194 10.0.0.195`
- On machine B define the stripe interface endpoint addresses and bring up the stripe, as follows:
Machine B > `ifconfig is0 10.0.0.195 10.0.0.194`

At this point you should now be able to verify the configuration of the stripe and test its connectivity.

Verifying a Network Stripe

This section describes how you can verify that your network stripe is correctly configured.

Procedure 10-3 Checking Stripe Configuration

The following steps describe how to use the `ifconfig(1M)` command to ensure that a stripe configuration is set-up correctly.

- Execute the `ifconfig` command, as follows:

```
> ifconfig -a
```

Output similar to the following appears:

```
is0: flags=8050d1<UP, POINTOPOINT, RUNNING, NOARP, CKSUM, DRVRLOCK, HIGHBW>  
    inet 10.0.0.194 --> 10.0.0.195 netmask 0xff000000  
    link local 192.168.2.194 remote 192.168.2.195  
    link local 192.168.1.194 remote 192.168.1.195
```

The output should show an additional entry for the striping interface (`is0`) similar to what is shown above. There should be a link line for each link pair you entered and the addresses displayed should match the addresses you entered in the configuration stage in the form:

```
link local laddr remote raddr
```

A verbose listing will also show the default send and receive socket space the interface will use if it has been configured. See “Tuning a Network Stripe Interface” on page 294, for more information on how to configure these parameters.

Procedure 10-4 Verifying Stripe Connectivity

Once it has been confirmed that the stripe interface configuration at both the local and remote ends are correct and both interfaces are up, you can test the stripe interface with `ping(1)`.

Using the remote stripe interface address of the 2-stripe configuration example above, you should see the output similar to the following if your striping software is working correctly:

```
> ping 10.0.0.195
PING 10.0.0.195 (10.0.0.195): 56 data bytes
64 bytes from 10.0.0.195: icmp_seq=0 ttl=254 time=0.488 ms
64 bytes from 10.0.0.195: icmp_seq=1 ttl=254 time=0.373 ms
64 bytes from 10.0.0.195: icmp_seq=2 ttl=254 time=0.335 ms
64 bytes from 10.0.0.195: icmp_seq=3 ttl=254 time=0.379 ms
64 bytes from 10.0.0.195: icmp_seq=4 ttl=254 time=0.327 ms

----10.0.0.195 PING Statistics----
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max = 0.327/0.380/0.488 ms
>
```

Monitoring both ends of the connection using the `netstat -C` command is also useful at this point to determine if pings have been sent down each of the physical links and if problems are occurring, to isolate where packets are going missing.

Once you have verified connectivity in this manner, you can run further tests using `ttcp(1)` or the test utility of your choice to exercise the interface. If you have problems, please refer to the Troubleshooting section below.

Stripe Interface MTUs

This section describes the optimal Maximum Transmission Unit (MTU) sizes to use with SGI Infinite Network Bandwidth software and covers the following topics:

- “Default MTU Calculation” on page 291
- “MTUs and Performance” on page 291

Default MTU Calculation

The MTU of a stripe interface is determined by the MTU of the physical links that the stripe uses. When a stripe is brought up, it scans each of the physical link interfaces and determines what the smallest interface MTU is from their current MTU settings. It then uses this value as the basis for the stripe interface MTU.

The stripe interface MTU is slightly smaller than the smallest MTU of the physical links as some space is taken up by headers the stripe adds to each packet. The overhead is automatically taken into account by the stripe interface when it calculates it's MTU.

Note that this MTU calculation means that it is not really worthwhile mixing physical links that have different MTUs as the stripe interface will not make use of the large MTUs that some links can provide and therefore does not benefit from them.

MTUs and Performance

Currently, the upper bound of performance for a stripe interface is determined by the number of packets per second that a stripe interface can deliver in order to the receiver. This is largely due to the fact that TCP will classify out of order packets as congestion and throttle throughput. Hence the stripe interface needs to deliver packets in order to TCP to maximize throughput.

As a result of this, the stripe interface needs to serialize the packet delivery at some point, and at this point the stripe is limited by the number of packets a single CPU can process. As of the IRIX 6.5.23 release, a stripe interface can reorder approximately 135-140,000 packets per second.

To put this number in context, 140,000 packets per second is approximately fifteen 100Mbit Ethernet interfaces full of traffic or 1.75 Gigabit Ethernet interfaces using 1500-byte MTUs.

When this limitation is reached, there are only two possible ways to improve throughput. The first way is to decrease the number of packets needing to be sent. In other words, the amount of data in each packet needs to be increased, in each packet and this requires interfaces to have larger MTUs. For Gigabit Ethernet, the solution is to use jumbo frames. The second way is to decrease the amount of CPU time required to process each packet.

By default, the reorder thread floats around between CPUs - it will tend to run on one of the NIC interrupt CPUs (usually bounce between them), but you can bind it to a separate CPU. For example, you could do this for a 2-link stripe:

- Bind reorder thread to CPU 1
- Bind link 1 NIC interrupts to CPU 2
- Bind link 2 NIC interrupts to CPU 3

Gigabit Ethernet with Jumbo Frames

For Gigabit Ethernet, large gains in throughput and decreases in CPU utilization can be attained by using jumbo (9000-byte) frames rather than standard (1500-byte) frames.

A link using jumbo frames running at 100% utilization will only send approximately 15,000 packets per second compared to 85,000 packets per second for a 1500-byte MTU link. Hence, the packet reordering limitation will not be an issue until a stripe with many more links is configured.

Using jumbo frames also reduces the CPU overhead in transmitting and receiving data across the link. Less work needs to be done to send the same amount of data, and hence we use a lot less CPU resources.

SGI recommends that you do **not** use jumbo frames for stripe interfaces over eg Gigabit Ethernet interfaces as there are hardware limitations that make them unreliable in striping environments.

Gigabit Ethernet with L2TCPSEG

If each of the local link interfaces is a `tg` Gigabit Ethernet interface and they all have the `IFF_L2TCPSEG` capability enabled (new in IRIX 6.5.23), the stripe interface will operate in a slightly different mode. Use `ifconfig(1)` to determine if all your interfaces have this capability enabled, as follows:

```
tg0: flags=8f15c43<UP, BROADCAST, RUNNING, FILTMULTI, MULTICAST, CKSUM, DRVRLOCK, LINK0, L2IPFRAG, L2TCPSEG, IPALIAS, HIGHBW, IPV6>
    inet 192.168.1.195 netmask 0xffffffff broadcast 192.168.1.255
    rcvspace 262144 sendspace 262144
    speed 1000.00 Mbit/s full-duplex
tg1: flags=8f15c43<UP, BROADCAST, RUNNING, FILTMULTI, MULTICAST, CKSUM, DRVRLOCK, LINK0, L2IPFRAG, L2TCPSEG, IPALIAS, HIGHBW, IPV6>
    inet 192.168.2.195 netmask 0xffffffff broadcast 192.168.2.255
```



```
recvspace 262144 sendspace 262144
speed 1000.00 Mbit/s full-duplex
```

The capability can be enabled for tg class interfaces by changing options in `/etc/config/tgconfig.options`. For specific information on how to enable this functionality, please consult the `tgconfig(1M)` man page.

When all the interfaces have this flag set, the stripe interface will set its MTU to 32992 bytes and the stripe interface will require less CPU time to send TCP or UDP data.

The receive side processes packets in exactly the same manner as previously, but due to the different internal fragmentation of the data being sent, the CPU overhead of reordering the packets is reduced by approximately 10-15%. With this mode enabled, a stripe interface can reorder approximately 150-155,000 packets per second.

Note: Even if stripe endpoints have different MTUs (one end of the stripe link supports L2TCPSEG and the other does not), the stripe will still communicate correctly as long as the physical link pair MTUs are the same. This MTU change is transparent to the physical links and only affects the internal fragmentation of the data on the wire.

Note: It is highly recommended that if you use the L2TCPSEG mode of the stripe that you enable the strict IP fragment reassembly option on both machines as this provides significant improvement in performance when packets are dropped by the network. This should be done even if none of the traffic over the stripe interface is using UDP. This can be enabled by executing the following `sysctl(1)` command, as follows: `> sysctl ipv4_strict_reassembly 1`

Network Stripe Monitoring

Monitoring of a network stripe interface may be done through all of the existing network interface monitoring tools such as `netstat(1)` and PCP. The kernel driver for the stripe interfaces also exports a number of metrics to PCP to indicate the health and efficiency of the stripe interfaces on the machine.

The metrics exported from the kernel are global metrics. That is, the statistics are accumulated from all active stripe interfaces and are not broken out into per interface

statistics. Despite this, a stripe that is not operating correctly can still be diagnosed from the exported statistics as error conditions tend to be uncommon. For help in diagnosing problems, you should refer to the long help text for each PCP metric available for the stripe interface.

The stripe metrics are exported from the `network.is` node. You can get a list of the metrics and a one line description of each metric from `pminfo(1)` by executing the command, as follows:

```
> pminfo -t network.is
```

A more detailed description of each metric can be obtained by from `pminfo(1)` by executing the command, as follows:

```
> pminfo -T network.is
```

The current values of each metric can be obtained from `pminfo(1)` by executing the command, as follows:

```
> pminfo -f network.is
```

An far more useful way of viewing the stripe metrics is to use `pmchart(1)`. This provides graphical views of the metrics over time and is the preferred method of real-time viewing of the metric values.

The `pmchart` software is part of the licensed PCP monitoring product. For more information on `pmchart`, see Chapter 4. “Monitoring System Performance” in the *Performance Co-Pilot for IRIX Advanced User’s and Administrator’s Guide*.

For more information on how to use `pminfo(1)` and `pmchart(1)`, see their respective man pages.

Tuning a Network Stripe Interface

This section describes how to tune a network stripe interface and covers the following topics:

- “Stripe Interface System Tunable Parameters” on page 295
- “ifconfig Parameters” on page 295
- “Binding and Placement” on page 297

- “tg Gigabit Ethernet Interfaces” on page 297
- “The Reorder Thread” on page 298
- “Applications” on page 298

Stripe Interface System Tunable Parameters

There are two global stripe system tunable parameters that can affect performance of a stripe interface. These are, as follows:

- `is_rrorder` - default value = 5

This controls the size of the send quota for each physical link. It determines how many packets the stripe interface feeds to a link interface before moving onto the next link interface. The quota is $2^{\text{is_rrorder}}$.

The more links you have in your interfaces, the more likely you are to want a smaller `is_rrorder` so you can keep packets on every link. The alternative to reducing `is_rrorder` is to increase the default send and receive socket buffer sizes on the stripe interface (see `ifconfig` Parameters below). This allows more packets to be sent at once and hence you are more likely to be able to keep each link busy.

- `is_reorder_flush_delay` - default value = 5

This is the delay in ticks (1 tick = 10ms) between executions of the window stall detection routine. This determines how quickly lost packets are detected.

On links that suffer from frequent packet loss, setting a lower value may help maintain optimal throughput. It is not generally advisable to increase this parameter as it can cause undesirable interactions with TCP timeouts and retries that will further harm throughput when packet loss occurs.

`ifconfig` Parameters

The following `ifconfig` parameters may need to be set to allow applications using a stripe interface to perform well with application specific tuning:

- `[-]highbw`

This needs to be set when the stripe output packet rate goes over 30-40,000 packets per second. This reduces the CPU overhead of TCP by reducing the number of ACKs that are sent.

Note that close attention must be paid if the stripe interface is making use of the L2TCPSEG functionality, as the stripe interface output packet rate may be substantially lower than the physical link output packet rate.

- `sspace n`

This is the default send socket buffer space in bytes that applications will be given when they create a TCP socket that uses the stripe interface. If you do not specify a value, the system default (60KB) will be used.

In general, the higher the bandwidth of the stripe interface, the larger *n* needs to be. Experimentation is the only way to determine the best value for a given configuration. However, the following table can be used as a basis for experimentation.

Table 10-1 Throughput and Estimated Socket Buffer Size

Throughput	Estimated Socket Buffer Size
less than 50MB/s	System default (60KB) to 128KB
50-100MB/s	128KB to 256KB
100-150MB/s	256KB to 512KB
150-200MB/s	512KB to 1MB
greater than 200MB/s	1MB to 2MB

- `rspace n`

This is the default receive socket buffer space in bytes that applications will be given when they create a TCP socket that uses the stripe interface. If you do not specify a value, the system default (60KB) will be used.

In general, tuning of this parameter is identical to `sspace`, and more often than not, will have the same value as `sspace`.

For further information on the configuration of these parameters, please refer to the `ifconfig(1M)` man page.

Binding and Placement

In general, binding of physical network interface interrupts and placement of applications becomes more important as the bandwidth of a stripe interface increases. As bandwidth increases, the latencies inherent in nonuniform memory access (NUMA) architectures will play a part in the maximum throughput that an application can achieve.

For example, it would not be unexpected to see a 10% drop in maximum throughput for each hop between the node the application is running on and the node on which the physical interface interrupt routines are running.

Hence, for high bandwidth stripe interfaces, it is desirable to cluster applications, physical interface interrupt routines and stripe reorder threads on CPUs and nodes as topologically close to each other as possible. This is usually necessary for stripe interfaces that make use of Gigabit Ethernet physical interfaces.

For low bandwidth stripe interfaces, there is no noticeable advantage to carefully clustering applications, interrupt routines and reorder threads close together - the bandwidth is not high enough for memory latency differences to affect throughput.

Note that the definition of a low bandwidth stripe interface will change with CPU power and machine architecture; a low bandwidth stripe on a Origin 3000 class machine with 700MHz R14000 CPUs is likely to be considered a high bandwidth stripe on a Origin 2000 with 195Mhz R10000 CPUs. Hence, the binding and placement considerations for the stripe interfaces on these two machines are going to be very different despite the fact they will carry the same amount of data.

`tg` Gigabit Ethernet Interfaces

For `tg` class Gigabit Ethernet interfaces using standard frames (1500-byte MTU), it is recommended that you bind each physical interface to a separate CPU due to the interrupt load of these interfaces.

If you are running jumbo frames (9000-byte MTU), then you can bind multiple interfaces to the one CPU without seeing any performance degradations. How many you can bind to a single CPU will be dependent on your specific hardware.

For specific information on how to bind the interrupt CPU for a `tg` class Gigabit Ethernet interface using the `/etc/config/tgconfig.options` file, see the `tgconfig(1M)` man page.

The Reorder Thread

Several XTHREADS (for more information on thread classes, see the `realtime(5)` man page) are created when a stripe is configured. The only XTHREAD that will greatly effect performance is the `reorder` thread. The reorder thread of a stripe is responsible for delivering reordered packets to the upper network stack layers, and hence there are situations where the placement of this thread may improve performance.

When a new stripe interface `isX` is configured, a reorder thread called `is_reorderX` gets created. This thread can be bound to a particular CPU by editing the `/var/sysgen/system/irix.sm` file and adding the statement, as follows:

```
XTHREAD is_reorderX CPU Y
```

Where:

- `X` is the stripe interface unit number (`isX`) to which you wish the XTHREAD command to apply to; and,
- `Y` is the number of the CPU you to which you wish to bind the reorder thread to.

If you have multiple stripe interfaces, then you can add multiple statements to bind the different threads as required.

In normal situations, the reorder thread will move between CPUs running on whatever CPU the scheduler decides is most appropriate. Binding this thread to a particular CPU is useful, if you wish to tightly control the locality of applications using the stripe. In most cases, it is not necessary to bind the reorder thread to a CPU. However, binding the reorder thread to a particular CPU may improve stripe performance in some cases. It is best to bind this thread (which will use up to an entire CPU) on the same or neighbor node to the node on which your applications using the stripe interface are running.

Applications

The placement of applications using a stripe interface will affect the maximum bandwidth that the stripe interface can provide to the application. Ideally, an application

that requires as much of the bandwidth a stripe interface can provide should be located on the same node as the physical interface interrupts and the reorder thread.

However, this is not always possible given the CPU load that the physical interface interrupts and the reorder thread can generate. Hence placing the application on an adjacent node or splitting the interrupt and reorder thread load over two nodes and running the application on one (or both) of these nodes may be the best solution.

Placement of applications can be controlled in various ways. The simplest is to use `runon(1)` or to build CPU sets to define the locality in which the application can run. For more information of CPU sets, please refer to the `cpuset(5)` man page.

Application I/O Sizes

Given the bandwidth a stripe interface can provide, application I/O patterns can have major effect on throughput. Throughput can be limited on either the send or receive side if an application does not use optimal I/O patterns.

For optimal throughput, SGI recommends the following:

- All I/O buffers are page aligned and an exact multiple of page size. This allows the network stack to do zero-copy transmit and receive where possible. This significantly reduces CPU overhead of sending and receiving data.
- Transmit I/O buffers should not be reused or freed immediately after they are sent if you are using TCP. This will defeat the transmit zero-copy path and result in significant extra CPU load. Release or reuse the buffers after a short period of time (say 1-2 milliseconds) to ensure that the transmit path has finished with the buffers first.
- I/O buffers should be as large as possible to reduce the system call overhead as much as possible. That is, we have more time available to send or receive data because the percentage of time spent doing system calls is lower when we have larger buffers.

In general, buffer sizes in the range of a quarter to twice the socket buffer size are most likely to be optimal for high bandwidth applications. For example, if you have a stripe interface with a default socket buffer space of 1MB on both the send and receive sides, then it is likely that the optimal I/O buffer sizes will fall somewhere in the range of 256KB to 2MB.

It is worth noting that the optimal I/O buffer sizes will change from application to application, and will also vary between different stripe interface configurations. Experimentation is the only way to find the optimal buffer sizes for any given application and setup.

It is also worth noting that if the application has to do a non-trivial amount of data processing between sending or receiving data, the application may become CPU bound and will not be able to sustain the throughput a stripe interface can provide. In these cases, you can either run multiple instances of the application to until the stripe interface is saturated (for example, running multiple parallel FTP sessions), or the application may need to be threaded so it can make use of multiple CPUs and hence be able to send, receive and process the data without being limited by a single CPU.

Troubleshooting SGI Infinite Network Bandwidth Software

This section describes network stripe troubleshooting and covers the following topics:

- “Testing Connectivity” on page 300
- “Diagnosing Performance Problems” on page 302

Testing Connectivity

If any of the physical link pairs show problems communicating with each other, then the stripe will have trouble communicating. The first step in diagnosing any stripe problem is determining if there is a problem with a link pair. If there is a problem with a link pair, it should either be fixed or removed from the stripe interfaces before attempting to reuse the stripe interface. A basic procedure for discovering if there is a problem with a link pair is, as follows:

- Execute a ping test, as described in “Verifying a Network Stripe” on page 289, first across each physical link in the stripe set, and then across the stripe interface itself. This does the following:
 - Demonstrates that each physical link is configured correctly.
 - Ensures that each physical link pair has connectivity; and
 - If all links have connectivity, then the stripe interfaces should be able to communicate with each other.

- As root, execute a short flood ping test, first across each physical link in the stripe set, and then across the stripe interface itself. A simple flood ping test of 1024 packets is, as follows:

```
> ping -c 1024 -f 10.0.0.195
PING 10.0.0.195 (10.0.0.195): 56 data bytes
.

----10.0.0.195 PING Statistics----
1024 packets transmitted, 1024 packets received, 0.0% packet loss
round-trip min/avg/max = 0.208/0.215/0.686 ms
4456.7 packets/sec sent, 4465.9 packets/sec received
>
```

This does the following:

- Demonstrates that higher single packet rates do not trigger problems on the physical link pairs.
- Cycles the stripe interface through all of its physical links, hence potentially locating the link that may be causing problems.
- Execute a TCP throughput test using `ttcp`. Do not use the default parameters of `ttcp` because this will execute a test that is far too short to be useful. The test should run for at least 10 seconds to allow TCP to reach full speed. While the TCP test is running you should monitor the stripe interfaces for abnormalities as described “Diagnosing Performance Problems” on page 302.

If a problem is not discovered with a link pair, then the stripe interfaces at both ends should be brought down and then back up, and you should re-run the stripe interface ping tests.

If the problem only shows up in the TCP tests, then it may be your tuning of the stripe interfaces (for example, placement of threads, applications, and so on) that is non-optimal and is therefore contributing to reordering or efficiency degradations.

If the problem still persists, you should destroy both stripe interfaces and recreate them from scratch.

Diagnosing Performance Problems

If your application is not performing as you expect over a stripe interface, you should attempt to tune the stripe and/or applications according to the guidelines given in “Tuning a Network Stripe Interface” on page 294.

However, factors outside of application and stripe tuning can affect the performance of the stripe. Incorrect tuning can also lead to performance problems. The key to diagnosing these problems is the monitoring of the stripe driver PCP metrics as described in “Network Stripe Monitoring” on page 293.

The following metrics are indicators of packet loss or extreme packet reordering problems:

- `network.is.in_underflow`
- `network.is.drain_underflow`
- `network.is.window_stalls`
- `network.is.window_flush_null`
- `network.is.window_flush_skipped`
- `network.is.err_bad_version`

The following metrics are indicators of packet reordering problems:

- `network.is.up_disordered`
- `network.is.window_stalls`
- `network.is.window_seqno_fixup`
- `network.is.window_flush_nlinks`

The following metrics can be used to gain insight into the efficiency of your stripe configuration:

- `network.is.in_window`
- `network.is.in_overlap`
- `network.is.up_ordered`
- `network.is.link_empty_headers`
- `network.is.link_header_alloc`

- `network.is.link_soft_cksums`
- ratio of `network.is.outq_wakeups` to `network.is.outq_drains`
The higher the ratio, the more efficient the stripe interface output.
- ratio of `network.is.reorder_wakeups` to `network.is.reorder_drains`
The higher the ratio, the more efficient the stripe interface reordered packet delivery.

Another common performance problem is that the stripe interface throughput varies between different speeds. This is commonly caused by the application being rescheduled to run on a node closer or further away from the stripe interface. If you see this happening, it is worthwhile monitoring what CPU your application is running on and determining if the change in throughput coincides with changes in the CPU the application is running on. If this is the problem, you should consider binding your application to a CPU or `cpuset` to stop this variation from occurring. See “Binding and Placement” on page 297, for details on how to do this.

BIND Standard Resource Record Format

The Berkeley Internet Name Domain (BIND) server uses a specific record format for the name server data files. This appendix details BIND's standard resource record format by resource record type in the following sections:

- “BIND Standard Resource Record Syntax” on page 306
- “About TTLs in BIND Resource Records” on page 306
- “About Special Characters in BIND Resource Records” on page 307
- “Specifying \$INCLUDE in BIND Resource Records” on page 307
- “Specifying \$ORIGIN in BIND Resource Records” on page 308
- “Specifying SOA—Start of Authority in BIND Resource Records” on page 308
- “Specifying NS—Name Server in BIND Resource Records” on page 309
- “Specifying A—Address in BIND Resource Records” on page 310
- “Specifying HINFO—Host Information in BIND Resource Records” on page 310
- “Specifying WKS—Well-Known Services in BIND Resource Records” on page 310
- “Specifying CNAME—Canonical Name in BIND Resource Records” on page 311
- “Specifying PTR—Domain Name Pointer in BIND Resource Records” on page 311
- “Specifying MB—Mailbox in BIND Resource Records” on page 311
- “Specifying MR—Mail Rename Name in BIND Resource Records” on page 312
- “Specifying MINFO—Mail Information in BIND Resource Records” on page 312
- “Specifying MG—Mail Group Member in BIND Resource Records” on page 312
- “Specifying MX—Mail Exchanger in BIND Resource Records” on page 312
- “Specifying RP—Responsible Person in BIND Resource Records” on page 313
- “Specifying TXT—Text in BIND Resource Records” on page 314

BIND Standard Resource Record Syntax

The records in the name server data files are called *resource records*. The Standard Resource Record (RR) Format is specified in RFC 1035. The standard format of resource records is

```
{name} {ttl} addr-class Record Type Record-specific data
```

- The first field is the name of the domain record. It must always start in column 1. For some RRs the name may be left blank, in which case it becomes the name of the previous RR.
- The second field is an optional time-to-live field, which specifies how long this data will be stored in the database. When this field is blank, the default time-to-live value is specified in the Start of Authority (SOA) resource record (described later in this section).
- The third field is the address class. Currently only the *IN* class (for Internet hosts and addresses) is recognized.
- The fourth field identifies the type of resource record.
- Subsequent fields depend on the type of RR.

Case is preserved in names and data fields when they are loaded into the name server. Comparisons and lookups in the name server database are not case sensitive.

About TTLs in BIND Resource Records

If you specify TTLs for resource records, it is important that you set them to appropriate values. The TTL is the amount of time (in seconds) that a resolver will use the data from your server before it asks your server again. If you set the value too low, your server will become loaded down with repeat requests. If you set it too high, information you change will not be distributed in a reasonable amount of time.

Most host information does not change much over time. A good way to set up your TTLs is to set them at a high value, and lower the value if you know a change is coming soon. You might set most TTLs between a day (86400) and a week (604800). When you know some data is changing soon, set the TTL for that RR to a low value, between an hour (3600) and a day, until the change takes place. Then reset it to its previous value. All resource records with the same name, class, and type should have the same TTL value.

About Special Characters in BIND Resource Records

The following characters have special meanings in resource records:

(blank)	A blank or tab character in the name field denotes the current domain.
@	A free-standing “at” sign (@) in the name field denotes the current origin.
.	A free-standing period in the name field represents the root domain name.
\x	The backslash designates that the special meaning of the character x does not apply. The x represents any character other than a digit (0–9). For example, use \. to place a dot character in a label.
\DDD	Each D is a digit; the complete string is the octet corresponding to the decimal number described by DDD. The octet is assumed to be text and is not checked for special meaning.
()	Parentheses enclose group data that crosses a line. In effect, newlines are not recognized within parentheses. This notation is useful with SOA and WKS records.
;	A semicolon precedes a comment; the remainder of the line is ignored.
*	An asterisk is a wildcard character.

Usually a resource record has the current origin appended to the name if the name is not terminated by a period (.). This scheme is useful for appending the current domain name to the data, such as workstation names, but can cause problems if you do not want the name to be appended. If the name is not in the domain for which you are creating the data file, end the name with a period. However, do not append the period to Internet addresses.

Specifying \$INCLUDE in BIND Resource Records

An include line has \$INCLUDE starting in column 1 and is followed by a filename. This feature helps you use multiple files for different types of data. For example:

```
$INCLUDE /usr/etc/named.d/mailboxes
```

This line is a request to load the file `/usr/etc/named.d/mailboxes`. The `$INCLUDE` command does not cause data to be loaded into a different zone or tree. It allows data for a given zone to be organized in separate files. For example, you might keep mailbox data separate from host data by using this mechanism.

Specifying \$ORIGIN in BIND Resource Records

`$ORIGIN` changes the origin in a data file. The line starts in column 1 and is followed by a domain origin. This feature is useful for putting more than one domain in a data file.

```
$ORIGIN Berkeley.EDU.
```

Specifying SOA—Start of Authority in BIND Resource Records

```
name {ttl}  addr-class SOA Source                Person-in-charge
@          IN          SOA ucbvax.Berkeley.EDU  kjd.ucbvax.Berkeley.EDU.
(
  1994021501;Serial
  10800    ;Refresh
  3600    ;Retry
  3600000 ;Expire
  86400   ;Minimum
)
```

The Start of Authority record, SOA, designates the start of a zone. There should be only one SOA record per zone.

The name is the name of the zone. It can be a complete domain name like `Berkeley.EDU`, or a name relative to the current `$ORIGIN`. The “at” sign (`@`) indicates the current zone name, taken from the “primary” line in the `named.boot` file or from a previous `$ORIGIN` line.

Source is the name of the host on which the master data file resides, typically the primary master server.

Person-in-charge is the mailing address for the person responsible for the name server. The mailing address is encoded in the form of a domain name where the “at” sign (`@`) separating the user name from the hostname is replaced with a period. In the example above, `kjd.ucbvax.berkeley.edu` is the encoded form of `kjd@ucbvax.berkeley.edu`.

Serial is the version number of this data file, and should be incremented whenever data are changed. Do not use floating point numbers (numbers with a decimal point, such as 1.1). A useful convention is to encode the current date in the serial number. For example, *25 April 1994 edit #1* is encoded as

```
1994042501
```

Increment the edit number if you modify the file more than once on a given day.

Refresh indicates how often, in seconds, a secondary name server is to check with the primary name server to see if an update is needed.

Retry indicates how long, in seconds, a secondary server is to retry after a failure to check for a refresh.

Expire is the maximum number of seconds that a secondary name server has to use the data before they expire for lack of getting a refresh.

Minimum is the default number of seconds to be used for the time-to-live field on resource records with no explicit time-to-live value.

Specifying NS—Name Server in BIND Resource Records

```
{name} {ttl}  addr-class  NS  Name server's_name
                IN           NS  ucbarpa.Berkeley.EDU.
```

The Name Server record, NS, lists the name of a machine that provides domain service for a particular domain. The name associated with the RR is the domain name, and the data portion is the name of a host that provides the service. Workstations providing name service need not be located in the named domain. There should be one NS record for each master server (primary or secondary) for the domain. If you use more than approximately 10 to 15 NS records for a zone, you may exceed DNS datagram size limits.

NS records for a domain must exist in both the zone that delegates the domain and in the domain itself. If the name server host for a particular domain is itself inside the domain, then a glue record is needed. A glue record is an Address (A) record that specifies the address of the server. Glue records are needed only in the server delegating the domain, not in the domain itself. For example, if the name server for domain SRI.COM is KL.SRI.COM, then the NS and glue A records on the delegating server look like this:

```
SRI.COM.      IN  NS      KL.SRI.COM.
KL.SRI.COM.  IN  A        10.1.0.2
```

The administrators of the delegating and delegated domains should ensure that the NS and glue A records are consistent and remain so.

Specifying A—Address in BIND Resource Records

```
{name}  {ttl}  addr-class  A      address
ucbvax  IN      A          128.32.133.1
        IN      A          128.32.130.12
```

The Address record, A, lists the address for a given workstation. The name field is the workstation name, and the address is the network address. There should be one A record for each address of the workstation.

Specifying HINFO—Host Information in BIND Resource Records

```
{name}  {ttl}  addr-class  HINFO  Hardware      OS
        IN      HINFO    SGI-IRIS-INDY  IRIX
```

The Host Information resource record, HINFO, is for host-specific data. This record lists the hardware and operating system running at the listed host. Only a single space separates the hardware information and the operating-system information. To include a space in the workstation name, you must place quotation marks around the name. There should be one HINFO record for each host. See the file */usr/etc/named.d/README* for the current list of names for IRIS4D Series workstations and servers. To learn the names for other types of hardware and operating systems, refer to the most current “Assigned Numbers” RFC (RFC 1340 as of this writing).

Specifying WKS—Well-Known Services in BIND Resource Records

```
{name}  {ttl}  addr-class  WKS  address      protocol  services list
        IN      WKS    192.12.6.16  UDP      (who route
        IN      WKS    192.12.63.16 TCP      timed domain)
        IN      WKS    192.12.63.16 TCP      (echo telnet
        IN      WKS    192.12.63.16 TCP      chargen ftp
        IN      WKS    192.12.63.16 TCP      smtp time
        IN      WKS    192.12.63.16 TCP      domain bootp
        IN      WKS    192.12.63.16 TCP      finger sunrpc)
```

The Well-Known Services record, WKS, describes well-known services supported by a particular protocol at a specified address. The list of services and port numbers comes from the list of services specified in */etc/services*. There should be only one WKS record per protocol per address.

Specifying CNAME—Canonical Name in BIND Resource Records

```
aliases {ttl}  addr-class  CNAME  Canonical name
ucbmonet      IN             CNAME  monet
```

The Canonical Name resource record, CNAME, specifies an alias for the official, or canonical, hostname. This record should be the only one associated with the alias name. All other resource records should be associated with the canonical name, not with the alias. Any resource records that include a domain name as their value (such as NS or MX) should list the canonical name, not the alias.

Aliases are also useful when a host changes its name. In that case, it is usually a good idea to have a CNAME record so that people still using the old name get to the right place.

Specifying PTR—Domain Name Pointer in BIND Resource Records

```
name {ttl}  addr-class  PTR  real name
6.130      IN             PTR  monet.Berkeley.EDU.
```

A Domain Name Pointer record, PTR, allows special names to point to some other location in the domain. The example of a PTR record given here is used to set up reverse pointers for the special IN-ADDR.ARPA domain. PTR names should be unique to the zone. Note the period (.) appended to the real name to prevent *named* from appending the current domain name.

Specifying MB—Mailbox in BIND Resource Records

```
name {ttl}  addr-class  MB  Machine
ben       IN             MB  franklin.Berkeley.EDU.
```

The Mailbox record, MB, lists the workstation where a user receives mail. The *name* field is the user's login. The machine field lists the workstation to which mail is to be delivered. Mailbox names should be unique to the zone.

Specifying MR—Mail Rename Name in BIND Resource Records

```
name {ttl}  addr-class MR  corresponding_MB
Postmaster IN          MR  ben
```

The Mail Rename Name record, MR, lists aliases for a user. The *name* field lists the alias for the name listed in the last field, which should have a corresponding MB record.

Specifying MINFO—Mail Information in BIND Resource Records

```
name {ttl}  addr-class MINFO  requests      maintainer
BIND       IN          MINFO  BIND-REQUEST  kjd.Berkeley.EDU
```

The Mail Information record, MINFO, creates a mail group for a mailing list. This resource record is usually associated with a Mail Group (MG) record, but can be used with a Mailbox (MB) record. The *name* is the name of the mailbox. The *requests* field is where mail (such as requests to be added to a mail group) should be sent. The *maintainer* is a mailbox that should receive error messages. This arrangement is appropriate for mailing lists when errors in members' names should be reported to someone other than the sender.

Specifying MG—Mail Group Member in BIND Resource Records

```
{mail group name} {ttl}  addr-class MG  member name
                   IN          MG  Bloom
```

The Mail Group record, MG, lists members of a mail group. Here is an example for setting up a mailing list:

```
Bind  IN  MINFO  Bind-Request      kjd.Berkeley.EDU.
      IN  MG     Ralph.Berkeley.EDU.
      IN  MG     Zhou.Berkeley.EDU.
```

Specifying MX—Mail Exchanger in BIND Resource Records

```
                                preference mail
name {ttl}  addr-class MX  value      exchanger
Munnari.OZ.AU.  IN          MX  10      Seismo.CSS.GOV.
*.IL.          IN          MX  10      CUNYVM.CUNY.EDU.
```

The Mail Exchanger record, MX, specifies a workstation that can deliver mail to a workstation not directly connected to the network. In the first example given here, Seismo.CSS.GOV. is a mail gateway that can deliver mail to Munnari.OZ.AU. Other systems on the network cannot deliver mail directly to Munnari. The two systems, Seismo and Munnari, can have a private connection or use a different transport medium. The preference value is the order that a mailer should follow when there is more than one way to deliver mail to a single workstation. See RFC 974 for more detailed information.

You can use a wildcard name containing an asterisk (*) for mail routing with an MX record. Servers on the network can state that any mail to a given domain is to be routed through a relay. In the second example given here, all mail to hosts in the domain IL is routed through CUNYVM.CUNY.EDU. This routing is done by means of a wildcard MX resource record, which states that *.IL has an MX of CUNYVM.CUNY.EDU.

Specifying RP—Responsible Person in BIND Resource Records

```
owner {ttl} addr RP mbox_domain_name      TXT_domain_name
franklin IN RP franklin.berkeley.edu admin.berkeley.edu.
```

The Responsible Person record, RP, identifies the name or group name of the responsible person for a host. Often it is desirable to be able to identify the responsible entity for a particular host. Otherwise, when that host is down or malfunctioning, it is difficult to contact someone who can resolve the problem or repair the host.

The *mbox_domain_name* field is a domain name that specifies the mailbox for the responsible person. Its format in master files uses the DNS convention for mailbox encoding, identical to that used for the Person-in-charge mailbox field in the SOA record. In the example given here, the *mbox_domain_name* shows the encoding for ben@franklin.berkeley.edu. You can specify the root domain name (just ".") to indicate that no mailbox is available.

The last field is a domain name for which TXT RRs exist. You can perform a subsequent query to retrieve the associated TXT resource records at the TXT domain name. Retrieving the TXT records provides a level of indirection so that the entity can be referred to from multiple places in the DNS. You can specify the root domain name (just ".") for the TXT domain name to indicate that no associated TXT RR exists. In the example, sysadmins.berkeley.edu is the name of a TXT record that could contain some text with names and phone numbers.

The format of the RP record is class insensitive. Multiple RP records at a single name may be present in the database. They should have identical TTLs.

The RP record is experimental; not all DNS servers implement or recognize it.

Specifying TXT—Text in BIND Resource Records

```
text-name    {ttl}    addr-class   TXT    text-data
location          IN          TXT    "Berkeley, CA"
```

The Text record, TXT, is used to hold descriptive text. The semantics of the text depend on the domain where it is found.

IRIX sendmail Reference

This appendix provides information about trace flags for debugging the *sendmail* command. For information about *sendmail* command syntax and usage, see the *sendmail(8)* man page.

Specifying the sendmail Debugging Flags

Several debugging flags are built into sendmail. Each flag includes a number and a level. The number identifies the debugging flag. The level, which defaults to 1, dictates how much information prints. A low level causes minimal information to print; a high level causes more comprehensive information to print. In general, levels greater than 9 cause so much information to print that it is of limited value. Debugging flags use the following syntax:

-d *debug-list*

A debug list includes the flag number and the flag level, as shown in the following examples:

- Set flag 13 to level 1.
-d13
- Set flag 13 to level 3.
-d13.3
- Set flags 5 through 18 to level 1.
-d5-18
- Set flags 5 through 18 to level 4.
-d5-18.4

Many debugging flags are of little use to the average sendmail user. Some are occasionally useful for helping to track down obscure problems. See the sendmail manpage for a list of common debugging flags.

sendmail Debugging Traceflags

The following list includes the sendmail debugging flags..

```
#      $Id: TRACEFLAGS,v 8.37 2002/05/24 23:37:32 ca Exp $
0, 4   main.c          main    canonical name, UUCP node name, a.k.a.s
0, 15  main.c          main    print configuration
0, 44  util.c          printav print address of each string
0, 101 main.c          main    print version and exit
1      main.c          main    print from person
2      main.c          finis
3      conf.c          getla, shouldqueue
4      conf.c          enoughspace
5      clock.c         setevent, clrevent, tick
6      savemail.c      savemail, returntosender
7      queue.c         queue name
8      domain.c        getmxrr, getcanonname
9      daemon.c        getauthinfo    IDENT protocol
9      daemon.c        maphostname
10     deliver.c       deliver
11     deliver.c       openmailer, mailfile
12     parseaddr.c     remotename
13     deliver.c       sendall, sendenvelope
14     headers.c       commaize
15     daemon.c        getrequests
16     daemon.c        makeconnection
17     deliver.c       hostsignature
18     usersmtp.c      reply, smtpmessage, smtpinit smtpmailfrom,
smtpdata
19     srvrsmtp.c      smtp
20     parseaddr.c     parseaddr
21     parseaddr.c     rewrite
22     parseaddr.c     prescan
24     parseaddr.c     buildaddr, allocaddr
25     recipient.c     sendtolist
26     recipient.c     recipient
27     alias.c         alias
27     alias.c         readaliases
27     alias.c         forward
```



```

27     recipient.c    include
28     udb.c          udbexpand, udbsender
29     parseaddr.c   maplocaluser
29     recipient.c   recipient (local users), finduser
30     collect.c     collect
30     collect.c     eatfrom
31     headers.c     chompheader
32     headers.c     eatheader
33     headers.c     crackaddr
34     headers.c     putheader
35     macro.c       expand, define
36     stab.c        stab
37     readcf.c      (many)
38     map.c         initmaps, setupmaps (bogus map)
39     map.c         map_rewrite
40     queue.c       queueup, orderq, dowork
41     queue.c       orderq
42     mci.c         mci_get
43     mime.c        mime8to7
44     recipient.c   writable
44     safefile.c    safefile, safedirpath, filechanged
45     envelope.c    setsender
46     envelope.c    openxsript
47     main.c        drop_privileges
48     parseaddr.c   rscheck
48     conf.c        validate_connection
49     conf.c        checkcompat
50     envelope.c    dropenvelope
51     queue.c       unlockqueue
52     main.c        disconnect
53     util.c        xfclose
54     err.c         putoutmsg
55     conf.c        lockfile
56     mci.c         persistent host status
57     util.c        snprintf
58     bf.c          bf* routines
60     map.c
61     conf.c        sm_gethostbyname
62     multiple      file descriptor checking
63     queue.c       runqueue process watching
64     multiple      Milter
65     main.c        permission checks
66     srvrsmtp.c    conformance checks
69     queue.c       scheduling
#if _FFR_QUARANTINE

```

```
70     queue.c           quarantining
#endif /* _FFR_QUARANTINE */
80     content length
81     sun remote mode
91     mci.c             syslogging of MCI cache information
94,>99  srvrsmtp.c       cause commands to fail (for protocol testing)
95     srvrsmtp.c       AUTH
95     usersmtp.c       AUTH
96     tls.c            Activate SSL_CTX_set_info_callback()
97     srvrsmtp.c       Trace automode settings for I/O
98     *                timers
99     main.c           avoid backgrounding (no printed output)
```

Index

A

- adding stations (with BIND), 178
- administration, system
 - documentation, xxvi
- aliases database
 - building, 266
 - format of, 266
 - rebuilding, 267
 - testing, 268
 - troubleshooting, 268
- anonymous FTP, 92
- anonymous FTP, how to use, 31
- arp command, 120
- ASSERT error messages (UUCP), 247

B

- backend, database, 192
- bandwidth, reserving, 105
- BIND
 - adding domains, 178
 - adding new stations, 178
 - and caching-only servers, 160
 - and forwarders, 164
 - and network planning, 33
 - and nslookup, 181
 - and the /etc/hosts file, 33, 44
 - boot file, 162
 - caching-only server setup, 174
 - caching-only servers, 163

- client-server configurations, 158
- database files, 161
- database management, 177
- debugging, 179
- deleting stations, 178
- /etc/config/named.options file, 165
- forwarding server setup, 175
- forwarding servers, 159
- localhost.rev file, 165
- management scripts, 178
- master servers, 159
- named.rev file, 165
- organization of, 156
- primary master server, 163
- primary server setup, 169, 177
- root.cache file, 165
- secondary master server, 163
- secondary server setup, 173
- server configurations, 158
- setup example, 168
- slave mode, 164
- slave servers, 159
- SYSLOG messages, 180

- BIND client resolve order, 193
- BIND database and MX records, 279
- BIND host lookup routines, 156
- boot file, BIND, 162
- bridge, definition of, 8

C

- cache, 188
- cache file format, 188
- cache files under UNS, 188
- cache tuning, 203
- caching-only server, BIND, 174
- caching-only servers, 160, 163
- chkconfig command, 61, 118
- client setup, BIND, 177
- configuring
 - network controllers, 4
 - network routers, 47
- connecting to the Ethernet, 40
- controller boards, and network troubleshooting, 130
- controller interface names, 4
- controllers (network), configuring, 4
- cu command (UUCP), 208

D

- daemon nsd, 186
- database backend files, 192
- database files, BIND, 161
- database management, BIND, 177
- DBM database, 267
- default lookup order, 189
- definition of network striping, 283
- deleting stations (with BIND), 178
- demand-dialing, 151
- Devices database (UUCP), 210, 211
- DHCP
 - client configuration, 85
 - network configuration, 68
 - relay-agent configuration, 83
 - server configuration, 70

- Dialcodes database (UUCP), 211, 223
- Dialers database (UUCP), 210, 211, 216
- dial-in configuration
 - PPP, 146
- dialing on demand, SLIP and PPP, 151
- dial-out configuration example
 - PPP, 142
 - SLIP, 142
- dial-out configuration with SLIP and PPP, 138
- DNS client resolve order, 193
- DNS protocol library, 199
- domain attributes, 201
- domain space, definition of, 156
- domains
 - definition of, 156
 - top level, 157
- domains (BIND), adding, 178

E

- editing the /etc/hosts file, 45
- electronic mail and network planning, 37
- /etc/alias.db file, 266
- /etc/chkconfi file, 43
- /etc/config/ifconfig.options file, 62
- /etc/config/ipaliases file, 61
- /etc/config/named.options file, 165
- /etc/config/netif.options file, 57, 58
- /etc/fstab file, 118
- /etc/hosts file
 - alternatives to, 33
 - editing, 45
 - host names in, 33, 44
 - router entries, 48
- /etc/hosts file, 118
- /etc/init.d/mail script, 261

/etc/init.d/network initialization script, 117
 /etc/init.d/network script, 47, 91, 119
 /etc/init.d/network.local script, 91
 /etc/nsswitch.conf file, 186, 189
 /etc/passwd/etc/passwd/usr/etc/remoteslip/etc/
 passwd, 144
 /etc/resolve.conf file, 189
 /etc/sys_id file, 45, 118
 Ethernet connections to an IRIS system, 40
 Ethernet error messages, finding, 42
 Ethernet, testing, 108-112

F

file attributes, 193, 202
 file transfer
 over SLIP and PPP, 152
 with FTP, 31
 filesystem namespace, 188
 .forward mail file, 278
 forwarders and BIND, 164
 forwarding server, BIND, 175
 forwarding, controlling on routers, 49
 FTP, anonymous, 92
 FTP, password protected, 97
 FTP, setting up account, 92
 FTP, using, 31
 Further, 191

G

gateway, definition of, 8
 genperm command (UUCP), 209

H

hardware options (network), 3
 hardware requirements (network), 1
 hostname resolution, 32
 hosts database
 modifying, 43
 router entries in, 47
 hosts database alternatives, 33

I

ifconfig command, 60, 119
 ifconfig-hy.options file, 118
 ifconfig.options file, 62
 inetd daemon, 118
 Infinite Network Bandwidth Overview, 282
 Infinite Network Bandwidth, definition, 283
 initializing the network, 117
 Integrated Services Digital Network (ISDN), 13
 Internet
 Gateway, 13
 Gateway access, 104
 Internet addresses
 and BIND, 14
 and NIS, 14
 and subnetworks, 34
 classes of, 34
 obtaining, 17
 planning, 14
 Internet Control Message Protocol (ICMP), 46
 interpreting network statistics, 123
 IP address allocation, 36
 IP aliasing, 60
 IPaliases.options file, 61
 ipfiltered daemon, 118
 IRIS InSight servers, 99

IRIS InSight, using NFS, 99
IRIX administration
 documentation, xxvi

K

kernel parameters and networking, 132

L

LDAP (Lightweight Directory Access Protocol), 194
libraries, protocol, 197
library attributes, 202
link pair, 282
link-pair, 282
.local file, 188
local station (UUCP), configuring, 236
local stations (UUCP)
 identifying, 235
localhost entry in /etc/hosts file, 44
localhost.rev file, 165
logging remote access, 91
logical or virtual interface, 282
lookup order, 189

M

mail queue, 276
mail system components, 254
mail.aliases nsd map, 267
mapping, name-to-address, 32
maps, 192
Maxuuscheds file (UUCP), 232
Maxuuxqts file (UUCP), 232
MDBM files, 192

MDBM protocol library, 199, 200
mdbm_dump command, 197
modem connections, 3
modem selection, 137
modifying the hosts database, 44
mqueue directory, 264
MTU Discovery, 128
multicast and NIS, 54
multicast routing, 50
multicast tunneling, 53
multi-homed hosts, 62
multiple network interfaces, configuring, 67
MX records, 279

N

named daemon, 178
named server, definition of, 155
named.boot file, 162, 178
named.hosts file, 164
named.rev file, 165
namespace format, UNS, 195
namespace, filesystem, 188
name-to-address mapping, 32
naming stations, 46
netif.options file, 57, 58, 118
netmask
 setting, 55
netmask option, 55
Netscape Mail, 37
netsnoop command, 122
netstat command, 120
netstat command example, 126
NetVisualyzer, 123
network address masks, 55

- network applications planning, 37
- network connections, testing, 46
- network controllers, choosing, 10
- network daemons
 - checking status of, 43
- network equipment, planning, 8
- Network File System (NFS), 152
- network hardware
 - options, 3
 - requirements, 1
- Network Information Center (NIC), 17, 157
- network initialization, 117
- network interfaces
 - checking, 119
 - displaying, 120
 - modifying addresses of, 59
 - modifying names of, 58
 - multiple, 67
 - variables, 57
- network management
 - functions, 116
 - tools, 119
- network masks, setting, 55
- network master script, 117
- network parameters, changing, 62
- network performance factors, 129
- network planning
 - BIND, 33
 - device number, 10
 - electronic mail, 37
 - Internet addresses, 14
 - media selection, 9
 - network applications, 37
 - NFS, 37
 - NIS, 33
 - PPP, 13
 - security, 37
 - size considerations, 7
 - SLIP, 12
 - subnetworks, 34
 - traffic, 10
 - UUCP, 13
- network script, 117
- network scripts, creating, 91
- network shutdown, 117, 119
- Network SLIP connections (NSLIP), 149
- network software checkout, 43
- network software options, 5
- network software, IRIS standard distribution, 4
- network startup, 117
- network statistics, interpreting, 123
- Network Stripe connectivity, 300
- Network stripe performance, 291
- network stripes, 283
- network striping hardware requirements, 283
- network striping terminology, 282
- network troubleshooting
 - configuration variables, 130
 - hardware, 129
 - media, 130
 - packet size, 131
 - servers, 131
- networking, preparing an IRIS system for, 39
- NFS
 - and network planning, 37
- NFS and UNS, 193
- NIS
 - and multicasting, 54
 - and network planning, 33
 - and the /etc/hosts file, 33, 44
- NIS client resolve order, 190
- NIS database, 192
- NIS maps, 192
- NIS protocol library, 199
- NIS server setup, 197

- NIS with UNS, 189
- nisserv protocol library, 200
 - /ns namespace, 188
- nsd daemon
 - operation with NIS, 191
- nsd daemon, 187
- nslookup command, 181
- nsswitch.conf file, 187

P

- password-protected FTP, 97
- PC TCP/IP connections to an IRIX network, 132
- performance factors (network), 129
- Permissions database (UUCP), 211, 223
- physical interface, 282
- ping command, 41, 46, 121
- ping command example, 123
- planning
 - for network traffic, 10
 - Internet addresses, 14
 - network controllers, 10
 - network devices, 10
 - network equipment, 8
 - network performance, 9
 - network size, 8
- Point to Point Protocol (PPP)
 - and network planning, 13
 - dynamic addressing, 150
 - modem requirements, 137
 - routing, 146
- Point to Point Protocol(PPP), 134
- Poll database (UUCP), 231
- preparing for networking, 39
- primary IP address, 60
- primary master server, BIND, 163
- primary server, BIND, 169

- proclaim
 - client configuration, 85
 - description, 68
 - network configuration, 68
 - relay-agent configuration, 83
 - server configuration, 70
- protocol library
 - DNS, 199
 - files, 199
 - MDBM, 199, 200
 - NIS, 199
 - nisserv, 200
- protocol, definition of, 186

Q

- querying attributes, 202
- queue
 - mail
 - multiple, 259
- queue, mail, 276

R

- rarpd daemon, 118
- rebooting networked stations, 56
- remote access logging, 91
- remote stations (UUCP)
 - identifying, 235
- remote stations (UUCP), configuring, 240
- removing stations (with BIND), 178
- repeater, definition of, 8
- reserving bandwidth, 105
- resolv.conf file
 - example, 192
- resolv.conf file, 159
- resolve order

- BIND client, 193
- DNS client, 193
- in UNS, 186
- NIS client, 190
- resolver routines, purpose of, 155
- restricted FTP, 97
- root.cache file, 165
- route advertising, forcing, 62
- route command, 122
- router
 - configuring, 47
 - controlling forwarding on, 49
 - dual interfaces on, 47
 - multiple interfaces on, 48
- router, definition of, 8
- routing, SLIP and PPP, 146
- RPC server registration, 120
- rpcinfo command, 120
- RSVP, 105
- rtquery command, 121
- rup command, 46
- rwhod daemon, 118

S

- secondary master server, BIND, 163
- secondary server, BIND, 173
- security
 - and network planning, 37
- selecting network media, 9
- sendmail
 - design features, 256
 - functions of, 258
 - general description, 255
 - software components, 259
- sendmail daemon, 259, 261, 265, 276
- sendmail scripts, 261
- sendmail.cf file
 - general description, 263
- sendmail.cf.auto file, 264
- sendmail.hf file, 264
- sendmail.mc file
 - general description, 263
- sendmail.st file, 264
- Serial Line Interface Protocol (SLIP)
 - and bidirectional links, 150
 - and data compression, 134
 - and file transfers, 152
 - and NFS, 152
 - debugging, 152
 - modem requirements, 137
 - modems for, 137
 - modems settings for, 137
 - purpose of, 152
- Serial Line Internet Protocol (SLIP)
 - and network planning, 12
 - routing, 146
- serial line networks, 3
- setting
 - domain attributes, 201
 - file attributes, 202
 - library attributes, 202
 - network masks, 55
 - table attributes, 202
- setup
 - kernel, 132
- SGL, 283
- should, 262
- shutting down the network, 117, 119
- Signal Quality Error, 130
- slave mode, BIND, 164
- snmpd daemon, 118
- spray command, 121
- standalone mode, 118
- standard distribution network software, 4

- starting the network, 117
- station naming, 46
- stripe or stripe interface, 282
- subnetworks
 - and network masks, 55
 - and troubleshooting, 131
 - Internet addresses for, 34
 - planning, 34
- Sysfiles database (UUCP), 231
- SYSLOG, troubleshooting with, 203
- system administration
 - documentation, xxvi
- Systems database (UUCP), 211, 219

T

- table attributes, 202
- table, description of, 187
- TCP/IP connectivity to personal computers, 132
- TCP/IP tuning, 127
- testing connectivity, 46
- testing packet reception, 121
- The, 283
- There, 283
- This, 282, 289, 290
- timed daemon, 118
- timeslave daemon, 118
- To, 286
- traceroute command, 122
- transferring files with FTP, 31
- transmitters, and network troubleshooting, 130
- troubleshooting
 - network connectivity, 121
 - network hardware, 129
 - network performance, 122
 - network traffic, 122

- nsd, 203
- routing, 121
- RPC mismatches, 120
- SLIP and PPP, 152

- ttcp
 - example, 124
- ttcp command, 122
- tunable parameters and networking, 132
- tuning TCP/IP, 127
- tunnelling, 282
- tunnels, 53

U

- Unified Name Service, 185
- UNIX to UNIX Copy Program, see UUCP
- unknown file (UUCP), 233
- UNS
 - and NFS, 193
 - cache files, 188
 - configuration file, 196
 - overview, 186
 - with NIS, 189
- UNS namespace format, 195
- UNS protocol libraries, 197
 - /usr/bin/rup command, 122
 - /usr/etc/arp command, 120
 - /usr/etc/configmail script, 261
 - /usr/etc/ifconfig command, 119
 - /usr/etc/named.d/named.boot file, 162
 - /usr/etc/named.reload script, 178
 - /usr/etc/named.restart script, 179
 - /usr/etc/netstat command, 120
 - /usr/etc/ping command, 121
 - /usr/etc/resolv.conf, 166
 - /usr/etc/route command, 122

/usr/etc/rpcinfo command, 120
 /usr/etc/rtquery command, 121
 /usr/etc/spray command, 121
 /usr/etc/traceroute command, 122
 /usr/etc/ttcp command, 122
 /usr/lib/aliases database, 266
 /usr/lib/aliases file, 264
 /usr/lib/sendmail.cf file, 263
 /usr/lib/sendmail.cf_m4 directory, 265
 /usr/lib/sendmail.hf file, 264
 /usr/lib/sendmail.st file, 264
 /usr/local file, 193
 /usr/mail directory, 265
 /usr/spool/mqueue directory, 264
 uucheck command (UUCP), 209
 uucico daemon (UUCP), 209, 210
 uucleanup command (UUCP), 209
UUCP
 administrative commands, 208
 administrative files for, 233
 and direct links, 207
 and modem setup, 207
 and network planning, 13
 and TCP/IP, 206, 245
 and telephone lines, 207
 configuring local stations, 236
 configuring remote stations, 240
 daemons, 209
 definition of, 205
 error messages for, 247
 hardware requirements, 207
 identifying local stations, 235
 identifying remote stations, 235
 physical connections for, 236
 setting up, 235
 supporting databases, 210
 testing connections for, 243
 user commands, 207

uucp command (UUCP), 208
 uugetty program (UUCP), 210
 uulog command (UUCP), 209
 uupick command (UUCP), 208
 uustat command (UUCP), 208
 uuto command (UUCP), 208
 Uutry command (UUCP), 209
 uux command (UUCP), 208
 uuxqt daemon (UUCP), 210

V

/var/sysgen/master.d/bsd file, 132

Y

y, 132
 ypbind, historical function, 190
 ypinit
 changes, 191
 command, 197
 ypmake command, 197
 ypserv, historical function, 190

Z

zones, definition of, 158

