HP-UX Internet Services Administrator's Guide

HP-UX 11i v2, HP-UX 11i v3

HP Part Number: B2355-91060
Published: February 2007
Edition: 2

Legal Notices

Copyright Notice

Table of Contents

List of Figures

List of Tables

List of Examples

About This Document

This document provides an overview of the Internet Services software and describes how to install and configure it on your operating system.

It is one of the documents available for the Internet Services suite of products. For a list of other Internet Services documents, see "Related Documentation" (page 14). These documents replace the document *Installing and Administering Internet Services* (B2355-90685), which was shipped with releases prior to the HP-UX 11i v2 operating system.

New and Changed Information in This Edition

The following sections are added to this document:

- "Restructuring of the InternetSrvcs Product" (page 18)
- "The ramD Routing Daemon" (page 23)

Intended Audience

This manual is intended for system and network administrators responsible for configuring and maintaining the Internet Services software on the HP-UX 11i v2 or HP-UX 11i v3 operating system. Administrators are expected to have knowledge of operating system concepts, commands, and the various routing protocols. It is also helpful to have knowledge of Transmission Control Protocol/Internet Protocol (TCP/IP) networking concepts and network configuration; this manual is not a TCP/IP tutorial.

HP-UX Release Name and Release Identifier

Each HP-UX 11i release has an associated release name and release identifier. The *uname*(1) command with the `-r` option returns the release identifier. The following table lists the releases available for HP-UX 11i.

Release Identifier	Release Name	Supported Process Architecture
B.11.11	HP-UX 11i v1	PA-RISC
B.11.20	HP-UX 11i v1.5	Intel® Itanium® Processor Family
B.11.22	HP-UX 11i v1.6	Intel Itanium Processor Family PA-RISC
B.11.23	HP-UX 11i v2.0	Intel Itanium Processor Family PA-RISC
B.11.31	HP-UX 11i v3	Intel Itanium Processor Family PA-RISC

Publishing History

The following table lists the publishing details of this document for various HP-UX releases.

Document Manufacturing Part Number	Operating System Supported	Publication Date
B2355–90774	11i v2	August 2003
B2355–91060	11i v2, 11i v3	February 2007

What's In This Document

The *HP-UX Internet Services Administrator's Guide* is organized as follows:

Chapter 1 **Introduction to Internet Services** Provides an overview of the different products supported in the Internet Services product suite.

Chapter 2 **Installing and Configuring Internet Services** Describes how to configure the Internet Services software on your system.

Chapter 3 **TCP Wrappers** Provides an overview of TCP Wrappers and explains how to enable IPv6 support and logging.

Chapter 4 **Configuring NTP** Describes the basic and advanced NTP concepts, components and configuration instructions required to use NTP. This chapter also includes troubleshooting information.

Chapter 5 **Troubleshooting Internet Services** Describes how to troubleshoot the Internet Services software.

Related Documentation

For more information about the Internet Services suite of products, see the following documents:

- *HP-UX Mailing Services Administrator's Guide* at:

  `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`

- *HP-UX Routing Services Administrator's Guide* at:

  `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`

- *HP-UX IP Address and Client Management Administrator's Guide* at:

  `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`

- *HP-UX Remote Access Services Administrator's Guide* at:

  `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`

- *HP-UX ramD Administrator's Guide* at:

  `http://docs.hp.com/en/netcom.html#Routing`

- *Using HP-UX Internet Services* at:

  `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`

- Request for Comments (RFC) at:

  `http://www.ietf.org/rfc.html`

- Other Documents

  For detailed technical and conceptual information about BIND, as well as information about planning a BIND hierarchy and using Sendmail with BIND, HP recommends that you read Paul Albitz and Cricket Liu, 2001. *DNS and BIND*. O'Reilly and Associates, Inc., This book is available at:

  `http://www.ora.com`

- iknow Topics of Interest

  HP iknow Topics of Interest describe some networking concepts and tasks, as well as other topics. You can find these documents on the HP-UX networking communications home page at the following URL:

  `http://docs.hp.com/iknow`

Typographical Conventions

This document uses the following typographic conventions:

`audit`(5)	An HP-UX manpage. In this example, *audit* is the name and *5* is the section in the *HP-UX Reference*. On the web and on the Instant Information CD, it may be a link to the manpage itself. From the HP-UX command line, you can enter “`man audit`” or “`man 5 audit`” to view the manpage. See `man` (1).
Book Title	The title of a book. On the web and on the Instant Information CD, it may be a link to the book itself.
`ComputerOut`	Text displayed by the computer.
`Command`	A command name or qualified command phrase, daemon, file, or option name.
$	The system prompt for the Bourne, Korn, and POSIX shells.
#	The superuser prompt.
`daemon`	Courier font type indicates daemons, files, commands, manpages, and option names.
`Variable`	The name of a variable that you may replace in a command or function or information in a display that represents several possible values.
[ ]	The contents are optional in formats and command descriptions. If the contents are a list separated by `|`, you can choose one of the items.

{}	The contents are required in formats and command description. If the contents are a list separated by |, you must choose one of the items.
(Ctrl+A)	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the plus.
Bold	The defined use of an important word or phrase.
...	The preceding element can be repeated an arbitrary number of times.
|	Separates items in a list of choices.

HP Encourages Your Feedback

HP welcomes your comments concerning this document. We are committed to providing documentation that meets your needs. Send your comments or suggestions to:

netinfo_feedback@cup.hp.com

Include the document title, manufacturing part number, and any comment or error found in this document. Also, include what we did right, so we can incorporate it into other documents.

1 Internet Services Overview

The HP-UX Internet Services software, (formerly the ARPA Services suite of products) enables your HP system to carry out the following tasks:

- Transfer files.
- Log on to remote hosts.
- Execute commands remotely.
- Manage IP addresses and network clients.
- Perform all routing protocols.
- Exchange mail with remote hosts on the network.
- Locate and configure networked services in enterprise networks.
- Start network services such as `ftp`, `telnet`, `rlogin`, and `tftp` without overloading the system.
- Implement a powerful security mechanism for various services spawned by `inetd`, the Internet super daemon.

The link product, which comes with the core operating system, provides the hardware and software needed for communication by an HP workstation over an IEEE 802.3 network, an Ethernet local area network, or an X.25 packet switch network. NIS and NFS Services also require link software and can run concurrently on the same node with the Internet Services.

The information in this manual applies to all HP workstations unless noted otherwise.

IMPORTANT: System Administration Manager (SAM) is deprecated in HP-UX 11i v3. HP System Management Homepage (HP SMH) is the system administration tool for managing HP-UX. HP SMH provides systems management functionality, at-a-glance monitoring of system component health, and consolidated log viewing. HP SMH provides Graphical User Interface (GUI), Text User Interface (TUI), and Command Line Interface (CLI) for managing HP-UX. You can access these interfaces using the `/usr/sbin/smh` command.

When you run either the `/usr/sbin/sam` or `/usr/sbin/smh` command and if the DISPLAY environment variable is set, HP SMH opens in the default Web browser. If the DISPLAY environment variable is not set, HP SMH opens using its terminal interface.

Introduction to Internet Services

The HP-UX Internet Services software combines services developed by the University of California at Berkeley (UCB), Cornell University, Merit Network, Inc., Carnegie-Mellon University (CMU), Hewlett-Packard, Massachusetts Institute of Technology (MIT), Internet Software Consortium, and other public domain sources.

ARPA services include the set of services developed by UCB for the Advanced Research Projects Agency (ARPA): ftp and telnet. ARPA services are used to communicate with HP-UX, UNIX®, and non-UNIX systems.

Berkeley services include the set of services developed by UCB to implement UCB protocols: BIND, sendmail, finger, the rexec library, rcp, rlogin, remsh, ruptime, rwho, and rdist. Berkeley Services are used to communicate with HP-UX or other UNIX systems.

The Internet Services software also contains several other services: BOOTP, tftp, rbootd, NTP, and DDFA.

We recommend that you also read the following books for more detailed technical and conceptual information:

- For Internet Services, see *TCP/IP Network Administration* by Craig Hunt, published by O'Reilly and Associates Inc.
- For BIND, see *DNS and BIND*, by Paul Albitz and Cricket Liu, published by O'Reilly and Associates, Inc.
- For sendmail, see *Sendmail, 2nd Edition*, by Bryan Costales with Eric Allman and Neil Richert, published by O'Reilly and Associates, Inc. You also can visit the Web site for sendmail:

  http://www.sendmail.org

You can get information about the O'Reilly books (including retail outlets where you can buy them, and how to order them directly from O'Reilly) by visiting the O'Reilly Web site:

http://www.ora.com

Restructuring of the InternetSrvcs Product

In earlier releases of HP-UX (prior to HP-UX 11i v2), the Internet Services suite of products comprised InternetSrvcs and MailUtilities.

Starting with HP-UX 11i v3, the InternetSrvcs products are restructured into 11 products based on the product functionality. However, the MailUtilities product is not modified. The restructured InternetSrvcs products are classified as required, recommended, and optional. This new structure provides the flexibility to deselect the recommended and optional products while installing the HP-UX 11i v3 operating system, and to remove the recommended and optional products after installing the HP-UX 11i v3 operating system. As a result, you need to install patches only for those products that are installed on your system. This improves software manageability of the system and security, and also reduces system downtime.

Table 1-1 describes the 11 products into which the InternetSrvcs product is restructured and the MailUtilities product.

Table 1-1 The Internet Services Products

Product Name	Bundle Name	Category	Bundle description
Restructured Products in the InternetSrvcs Product			
DHCPv4	HPUX-DHCPv4	Recommended	This bundle contains the Dynamic Host Configuration Protocol (DHCP) daemon, the Bootstrap Protocol (BOOTP) server, clients, utilities, and sample configuration files.
DHCPv6	HPUX-DHCPv6	Recommended	This bundle contains the DHCP product for IPv6, the DHCPv6 server, clients, utilities, and sample configuration files.
FTP	HPUX-FTPServer	Recommended	The bundle contains the File Transfer Protocol (FTP) server, clients, utilities, and sample configuration files.
Gated-Mrouted	HPUXGatedMrouted	Recommended	This bundle contains the gated and mrouted server, utilities and sample configuration files.
InternetSrvcs	HPUXMinRuntime	Required	This bundle contains the basic services required to start the system.
MailUtilities	HPUXMinRuntime	Required	This bundle contains the server, client, utilities, and sample configuration files related to the mailing utilities.
NameService	HPUX-NameServer	Recommended	This bundle contains the named server, clients, utilities, and sample configuration files.
NTP	HPUX-NTP	Recommended	This bundle contains the xntpd server, clients, utilities, and sample configuration files.
RAMD	HPUX-RAMD	Recommended	This bundle contains ramd server, clients, utilities, and sample configuration files.
Sendmail	HPUX-MailServer	Required	This bundle contains the Sendmail server, clients, utilities, and sample configuration files.
SLP	HPUX-SLP	Optional	This bundle contains the slpd server, clients, APIs, utilities, and sample configuration files.
TCPWrappers	HPUX-TCPWRAP	Recommended	This bundle contains tcpd server, clients, APIs, utilities, and sample configuration file.
MailUtilities Product			
MailUtilities	HPUXMinRuntime	Required	This bundle contains the server, client, utilities, and sample configuration files related to the mailing utilities.

For more information on the products in Section , see “Software Descriptions” (page 20).

Software Versions

Table 1-2 lists the product versions that have been made available with this version of Internet Services on the HP-UX 11i v2 and HP-UX 11i v3 operating systems. The software versions listed in this table are public domain versions.

Table 1-2 Software Versions

Software	Version
FTP	2.6.1
Sendmail	8.13.3
BIND	9.3.2
`gated`	3.5.9
`mrouted`	3.8
TCP Wrappers	7.6.1

Software Descriptions

This chapter provides an overview of Internet Services. It discusses the following topics:

- "The ftp Service" (page 20)
- "The tftp Service" (page 21)
- "The telnet Command" (page 21)
- "R-Commands" (page 21)
- "The elm Utility" (page 22)
- "The mail and mailx Utilities" (page 22)
- "The Sendmail Utility" (page 22)
- "BIND" (page 22)
- "DHCP" (page 22)
- "The gated Service" (page 23)
- "The mrouted Service" (page 23)
- "The ramD Routing Daemon" (page 23)
- "Secure Internet Services" (page 24)

The `ftp` Service

The `ftp` (File Transfer Protocol) service copies files among hosts on the network that support Internet Services. It runs on the client host and supports ASCII, binary, and tenex file transfer protocol types. ASCII is the default type. Whenever `ftp` establishes a connection between two similar systems, it automatically switches to the binary type. For more information, type `man 1 ftp` or `man 1M ftpd at the HP-UX prompt.`

The `tftp` Service

The `tftp` (Trivial File Transfer Protocol) service, used with `bootp` to enable some diskless systems (such as the HP 700/X terminal), transfers files containing bootstrap code, fonts, or other configuration information. You must invoke the `tftpd` server via `inetd`. Type `man 1 tftp` or `man 1M tftpd` at the HP-UX prompt for more information.

The `telnet` Command

The `telnet` command allows you to log on to a remote host that supports Internet Services. Type `man 1 telnet` or `man 1M telnetd` at the HP-UX prompt for more information.

R-Commands

You can access any remote machine by using the following commands:

- `rexec` You can connect to a remote UNIX host on the network and execute specific commands using rexec. It uses a library routine and the rexecd daemon for this task. Type `man 3N rexec` or `man 1M rexecd` at the HP-UX prompt for more information.
- `rcp` You can copy files, directory subtrees or a combination of files and directory subtrees between UNIX hosts using rcp. Type man 1 rcp at the HP-UX prompt for more information.
- `rlogin` You can log on from one UNIX system to another network using `rlogin`. Type `man 1 rlogin` or `man 1M rlogind` at the HP-UX prompt for more information.
- `remsh`

  `remsh` works in the same way as `rexec`, except that it executes commands from a remote shell. Type `man 1 remsh` or `man 1M remshd` at the HP-UX prompt for more information.
- `ruptime` You can obtain information about specified UNIX nodes running the `rwhod` daemon by running the `ruptime` command. It is not supported over X.25 networks or over networks using the PPL (SLIP) product. Type `man 1 ruptime` or `man 1M rwhod` at the HP-UX prompt for more information.
- `rwho`

  You can obtain information about specified users logged in on local hosts on the network running the `rwhod` daemon by running the `rwho` command. `rwho` is not supported over X.25 networks or over networks using the PPL (SLIP) product. Type `man 1 rwho` or `man 1M rwhod` at the HP-UX prompt for more information.
- `rdistd` You can distribute and maintain identical copies of files across multiple hosts on the network by running the `rdist` command. Type `man 1 rdist` at the HP-UX prompt for more information.

See "Installing and Configuring Internet Services" (page 25) for information on installing and configuring the previous services.

See *HP-UX Remote Access Services Administrator's Guide*, at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, for complete information about these services.

The `elm` Utility

`elm`'s screen-oriented interface runs with Sendmail or with any other UNIX Mail Transport Agent and enables you to read and compose mail messages. It supports an industry-wide MIME standard for nontext mail messages, a special forms message and forms reply mechanism, and an easy-to-use alias system for individuals and groups. Type `man 1 elm` at the HP-UX prompt for more information.

The `mail` and `mailx` Utilities

`mailx`, an interactive message processing system, provides an easy and flexible environment for exchanging mail messages electronically. Type `man 1 mailx` at the HP-UX prompt for more information.

The Sendmail Utility

The Sendmail service works with your network mailers (for example, `elm` and `mailx`) to perform internetwork mail routing among UNIX and non-UNIX hosts on the network. It allows you to exchange mail messages with other hosts on the local area network via gateways. Type `man 1M sendmail` at the HP-UX prompt for more information.

See "Installing and Configuring Internet Services" (page 25) for information on installing and configuring Sendmail.

See the *HP-UX Mailing Services Administrator's Guide* for detailed information on Mailing Services at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`.

BIND

BIND (Berkeley Internet Name Domain) implements the Domain Name System (DNS). BIND is a distributed database service that resolves host names and enables internetwork mail. See the *HP-UX IP Address and Client Management Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, or type `man 1M named`. at the HP-UX prompt for more information.

DHCP

Dynamic Host Configuration Protocol (DHCP) is an extension of `bootp` that defines a protocol for passing configuration information to hosts on a network. It automatically allocates reusable network addresses and reduces the cost of managing IPv4 and IPv6 nodes in environments where administrators require more control over the allocation of addresses. See the *HP-UX IP Address and Client Management Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, or type `man 1M dhcpv6d` at the HP-UX prompt for more information.

The `gated` Service

The `gated` service determines routing over the Internet. See the *HP-UX Routing Services Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, or type `man 1M gated` at the HP-UX prompt for more information.

NTP

Network Time Protocol (NTP) maintains the local clock on an HP-UX workstation in agreement with Internet-standard time servers. See "Getting Started with NTP" (page 47), or type `man 1M xntpd` at the HP-UX prompt for more information.

TCP Wrappers

The Transmission Control Protocol (TCP) Wrappers product suite provides an enhanced security mechanism for services spawned by the Internet Services daemon, `inetd`. See "TCP Wrappers" (page 37) for more information.

The `mrouted` Service

The `mrouted` service implements the Distance-Vector Multicast Routing Protocol (DVMRP) for routing IP multicast datagrams. See the *HP-UX Routing Services Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, or type `man 1M mrouted` at the HP-UX prompt for more information.

The ramD Routing Daemon

The Route Administration Manager (ramD) is a routing daemon that handles multiple Internet Protocol version 6 (IPv6) routing protocols. You can configure the ramD Daemon (`ramd`) to perform all or any combination of the supported protocol functions. `ramd` handles the following routing protocols:

- Routing Information Protocol Next-Generation (RIPng)
- Border Gateway Protocol (BGP)
- Border Gateway Protocol (BGP)
- Intermediate System - Intermediate System for IPv6 (IS-IS)

Upon startup, `ramd` reads the HP-UX kernel routing table on the local machine. It maintains a complete routing table in the user space and synchronizes this table with the HP-UX kernel routing table. See the *HP-UX ramD Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, or type `man 1M ramD` at the HP-UX prompt for more information.

Secure Internet Services

Secure Internet Services (SIS) is an optionally enabled mechanism that incorporates Kerberos V5 Release 1.0 authentication and authorization for the following services: `ftp`, `rcp`, `remsh`, `rlogin`, and `telnet`. See *Using HP-UX Internet Services*, at the URL `http://www.docs.hp.com/hpux/netcom/index.html` for more information.

2 Installing and Configuring Internet Services

This chapter describes how to install and configure the Internet Services software on your system. It discusses the following topics:

- "Installing the Internet Services Software" (page 25)
- "Configuring the Internet Services Software" (page 25)

Installing the Internet Services Software

The Internet Services software is packaged along with the core HP-UX 11i v2 and HP-UX 11i v3 operating systems. Do not create or modify any system file while installing the operating system on your machine. The core operating system creates and modifies the necessary files on your system automatically.

Configuring the Internet Services Software

This chapter describes how to configure the Internet Services software on your system. It discusses the following topics:

- "Configuring the Name Service Switch" (page 25)
- "Configuring an Internet Address" (page 27)
- "Configuring inetd" (page 31)
- "Configuring Logging" (page 33)

Configuring the Name Service Switch

The name service switch determines where your system will search for information stored in the following files:

- `/etc/mail/aliases`
- AutoFS maps (like `/etc/auto_master` and `/etc/auto_home`)
- `/etc/group`
- `/etc/hosts`
- `/etc/netgroup`
- `/etc/networks`
- `/etc/passwd`
- `/etc/protocols`
- `/etc/publickey`
- `/etc/rpc`
- `/etc/services`

For all types of information except host information, you can configure your system to use NIS (one of the NFS services) or the local `/etc` file, in any order.

For host information, you can configure your system to use BIND (DNS), NIS, or the `/etc/hosts` file.

The default name service switch configuration is adequate for most installations, so you probably do not have to change it. The default configuration is explained in the section "Default Configuration" (page 26).

Also, for more information about the name service switch configuration files supplied in the `/etc` directory, including the syntax of the configuration file and customizing your configuration, see *Installing and Administering NFS Services*, available at the URL `http://www.docs.hp.com/hpux/netcom/index.html`, or type `man 4 nsswitch.conf` at the HP-UX prompt.

Hostname Fallback

The ability to consult more than one name service for host information is often called **hostname fallback**. The name service switch provides **client-side hostname fallback**, because it is incorporated into client-side programs (for example, `gethostbyname`), which request host information.

The Network Information Service (NIS), one of the NFS services, allows you to configure a **server-side hostname fallback**. This feature causes the NIS server to query BIND when it fails to find requested host information in its database. The NIS server then returns the host information to the client through NIS. This server-side hostname fallback is intended for use with clients such as PCs, which do not have the name service switch feature. HP recommends that you use the name service switch if possible, instead of the server-side hostname fallback provided by NIS. For more information about the NIS server-side hostname fallback, see *Installing and Administering NFS Services*, at the URL `http://www.docs.hp.com/hpux/netcom/index.html`.

NOTE: Configuring the name service switch is a separate task from configuring the name services themselves. You must also configure the name services before using them. The name service switch just determines which name services are queried and in what order.

HP recommends that you maintain at least a minimal `/etc/hosts` file that includes important addresses like gateways, diskless boot servers and root servers, and your host's own IP address. HP also recommends that you include the word `files` in the `hosts` line to help ensure a successful system boot using the `/etc/hosts` file when BIND and NIS are not available.

Default Configuration

If the `/etc/nsswitch.conf` file does not exist, or if the line for a particular type of information is absent or syntactically incorrect, the following default configuration is used:

```
passwd:    files nis group:     files nis hosts:     dns [NOTFOUND=return] nis [NOTFOUND=return] files
networks:  nis [NOTFOUND=return] files protocols: nis [NOTFOUND=return] files rpc:       nis [NOTFOUND=return] files
```

```
publickey: nis [NOTFOUND=return] files netgroup:  nis [NOTFOUND=return] files automount: files nis aliases:   files nis
 services   nis [NOTFOUND=return] files
```

If your `/etc/nsswitch.conf` file contains a syntactically correct line for a particular type of information, that line is used instead of the default.

Troubleshooting Using the `nsquery` Command

You must use the `nsquery` command to troubleshoot the name service switch. The `nsquery` command displays the name service switch configuration that is currently in use. Then, it displays the results of the query. To perform a `hosts`, `passwd`, or `group` lookup, issue the following command at the HP-UX prompt:

```
/usr/contrib/bin/nsquery lookup_type lookup_query
```

The *`lookup_type`* can be `hosts`, `passwd`, or `group`. The *`lookup_query`* can be a host name or an IP address, a user name or user ID, or a group name or group ID.

The following example uses `nsquery` to perform a lookup of the host name `brock`:

```
# /usr/contrib/bin/nsquery hosts brock  Using “nisplus [NOTFOUND=return] files” for the hosts policy.  Searching
nisplus for brock brock was NOTFOUND  Switch configuration: Terminates Search
```

As an optional third argument to `nsquery`, you can supply a name service switch configuration in double quotes, as in the following example:

```
# /usr/contrib/bin/nsquery passwd 30 "files nis"  Using "files nis" for the passwd policy.  Searching /etc/passwd for 30
 User name: www User ID: 30 Group ID:  1 Gecos: Home Directory: / Shell:  Switch configuration: Terminates Search
```

For more information, type `man 1 nsquery` at the HP-UX prompt.

Configuring an Internet Address

This section describes how to configure your host to find other hosts on the network, by host name or IP address. It discusses the following topics:

- “Choosing a Name Service” (page 27)
- “Editing the /etc/hosts File” (page 28)
- “Configuring a Route” (page 29)
- “Changing a Host’s IP Address” (page 30)

Choosing a Name Service

HP-UX provides ways to translate host names to IP addresses or IP addresses to host names:

- BIND (Berkeley Internet Name Domain), which is Berkeley’s implementation of the Domain Name System (DNS).
- The `/etc/hosts` file, a simple ASCII file that is searched sequentially.
- NIS (Network Information Service), one of the NFS services (previously called “Yellow Pages”).

By configuring the name service switch, you can use these name services in any order you choose. See “Configuring the Name Service Switch” (page 25).

If you have a large network, or if you need to connect to Internet hosts outside your local network, use BIND as your primary name service. When you use BIND, you administer a central database containing only the hosts on your local network, and you have access to the databases on all the other hosts on the Internet. See “Configuring and Administering the BIND Name Service” in the *HP-UX IP Address and Client Management Administrator’s Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services`, for instructions on configuring BIND.

If you have a large network and little need for Internet connectivity, you can use NIS as your primary name service. The NIS hosts database is administered centrally on one of your hosts. However, this database must contain the names and IP addresses of all the other hosts in your network.

Alternatively, you can use the `/etc/hosts` file as your primary name service. Each host in your network needs a copy of the `/etc/hosts` file containing the names and addresses of all the other hosts in your network. For information on the `/etc/hosts` file, see “Editing the /etc/hosts File” (page 28).

NOTE: If you choose to use BIND or NIS as your primary name service, you still need to configure a minimal `/etc/hosts` file so that your host can boot if BIND or NIS is not available.

Editing the /etc/hosts File

You can use any text editor to edit the `/etc/hosts` file, or you can use the HP System Management Homepage (HP SMH).

Follow these steps to edit the `/etc/hosts` file:

1. If no `/etc/hosts` file exists on your host, copy `/usr/newconfig/etc/hosts` to `/etc/hosts`, or use `ftp` to copy the `/etc/hosts` file to your host from another host on your network. Type `man 1 ftp` at the HP-UX prompt for more information.
2. Make sure your `/etc/hosts` file contains the following line:

```
127.0.0.1       localhost       loopback
```

3. Add your own host’s IP address, name, and aliases to the `/etc/hosts` file, as in the following example:

```
15.13.131.213      hpindlpk      romney
```

The first field is the IP address, the second is the official host name (as returned by the `hostname` command), and any remaining fields are aliases. Type `man 4 hosts` at the HP-UX prompt for more information.

4. If your host has more than one network interface installed, add a line to `/etc/hosts` for each interface. The `/etc/hosts` entries for your host will have the same official host name but different aliases and different IP addresses.
5. Add any other hosts to the `/etc/hosts` file that you need to reach. If you use a BIND or NIS server on a different host, add that host to your `/etc/hosts` file.

   If you have no default gateway configured, and you add a host that is not on your subnet, HP SMH will prompt you for the gateway. To stop the prompting, configure a default gateway.
6. If you are not using HP SMH, you must configure a gateway for each host that is not on your subnet. See "Configuring a Route" (page 29).
7. Make sure the `/etc/hosts` file is owned by user `root` and group `other`, and make sure the permissions are set to 0444 (`-r--r--r--`).

Configuring a Route

To configure a route from your system to other networks, complete the following steps:

1. If you use only one gateway to reach all systems on other parts of the network, configure a default gateway.

   You can use HP SMH to configure a default gateway, or if you are not using HP SMH, issue the following command:

   ```
   /usr/sbin/route add default gateway_address 1
   ```

   where *`gateway_address`* is the IP address of the gateway host.

   Then, set the following environment variables in the `/etc/rc.config.d/netconf` file:

   ```
   ROUTE_DESTINATION[0]="default" ROUTE_GATEWAY[0]="gateway_address" ROUTE_COUNT[0]="1"
   ```

   If the default gateway is your own host, set the `ROUTE_COUNT` variable to 0. Otherwise, set it to 1.
2. If your host is a gateway, configure the destination networks that can be reached from its network interfaces. Issue the following command for each network interface on your host:

   ```
   /usr/sbin/route add net destination IP_address
   ```

   where *`destination`* is a network address reachable by your host, and *`IP_address`* is the address of the network interface.

Then, create a new set of routing variables in the /etc/rc.config.d/netconf file for each network interface. Whenever you create a new set of variables, increment the number in square brackets, as in the following example:

```
ROUTE_DESTINATION[1]="15.13.131.0" ROUTE_GATEWAY[1]="15.13.131.213" ROUTE_COUNT[1]="0"
```

3. If you will not be using gated, configure routes to all the networks you need to reach. Type the following command for each network you need to reach from your host:

   ```
   /usr/sbin/route add net network_address gateway_address
   ```

   Then, create a new set of routing variables in the /etc/rc.config.d/netconf file for each new route. Whenever you create a new set of variables, increment the number in square brackets.

   ```
   ROUTE_DESTINATION[n]="network_address" ROUTE_GATEWAY[n]="gateway_address" ROUTE_COUNT[n]="1"
   ```

   If ROUTE_GATEWAY[*n*] is your own host, set ROUTE_COUNT[*n*] to 0. Otherwise, set it to 1.

4. Type the following command to verify the routes you have configured:

   ```
   /usr/bin/netstat -r
   ```

For more information on static routing, type man 1M route or man 7 routing at the HP-UX prompt.

If you have a large and complicated network, use gated for dynamic routing. See “Configuring gated” in the *HP-UX Routing Services Administrator’s Guide* at the URL http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services for more information.

Changing a Host’s IP Address

When you use HP SMH to change a host’s IP address, HP SMH does not perform all these steps. For example, HP SMH does not update BIND or NIS databases. To change a host’s IP address, complete the following steps:

1. Change the host’s IP address in the /etc/hosts file. See “Editing the /etc/hosts File” (page 28) .
2. Change the IP_ADDRESS[*n*] variable in the /etc/rc.config.d/netconf file to the new IP address.
3. If the host is on a network that uses BIND, change the host’s IP address in the data files of the authoritative name servers. See “Configuring and Administering the BIND Name Service” in the *HP-UX IP Address and Client Management Administrator’s Guide* at the URL

`http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services` for more information.

If the host is on a network that uses NIS, change its IP address in the `/etc/hosts` file on the NIS master server, and issue the following commands to regenerate the `hosts` database and push it out to the NIS slave servers:

```
cd var/yp /usr/ccs/bin/make hosts
```

4. If the host is moving to a different subnet, change the `ROUTE_DESTINATION`, `ROUTE_GATEWAY`, and `BROADCAST_ADDRESS[`*n*`]` variables in `/etc/rc.config.d/netconf`.

   If the host is moving to a network that uses a different subnet mask, change the `SUBNET_MASK[`*n*`]` variable in `/etc/rc.config.d/netconf`.
5. If the host is moving to a different network, you may have to configure new routes for it. See "Configuring a Route" (page 29).
6. If the host is on a network that uses `gated`, change its IP address on all the `gated` routers. See "Configuring gated" in the *HP-UX Routing Services Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services` for more information.
7. If the host is a BOOTP client, change its IP address in the `/etc/bootptab` file on the BOOTP server. If the host is a BOOTP server, and a BOOTP relay agent is configured to relay boot requests to the host, change the host's IP address in the `/etc/bootptab` file on the BOOTP relay agent. See "Configuring TFTP and BOOTP Servers" in the *HP-UX Remote Access Services Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services` for more information.
8. If the host is an NTP server, change its IP address in the `/etc/ntp.conf` file on NTP clients. If the host is an NTP client and is moving to another network, you might have to configure a different NTP server in its `/etc/ntp.conf` file. See "Configuring the Network Time Protocol (NTP)" in the *HP-UX IP Address and Client Management Administrator's Guide* at the URL `http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services` for more information.
9. Reboot the host.

Configuring inetd

The Internet daemon, `/usr/sbin/inetd`, is the master server for many of the Internet Services. The `inetd` daemon listens for connection requests for the services listed in its configuration file and spawns the appropriate server on receiving a request.

You must invoke `inetd` as part of the boot process, by running the following command at the HP-UX prompt:

```
# /sbin/init.d/inetd start
```

The `/etc/inetd.conf` file is the `inetd` configuration file, which lists the services that may be started by `inetd`. In addition to the configuration file, you can configure an optional security file called `/var/adm/inetd.sec` to restrict access to the services started by `inetd`.

This section provides instructions for completing the following tasks:

- "Editing the /etc/inetd.conf File" (page 32)
- "Editing the /var/adm/inetd.sec File" (page 32)

If you want to write your own service and tie it in to `inetd`, see the *BSD Sockets Interface Programmer's Guide,* at the URL `http://docs.hp.com/hpux/onlinedocs/B2355-90136/B2355-90136.html` for more information.

Editing the /etc/inetd.conf File

To edit the `inetd.conf` file using any text editor, complete the following steps:

1. Make sure `/etc/inetd.conf` contains the following lines. If any of the lines start with a number sign (#), remove the number sign to enable that particular service.

   ```
   ftp    stream tcp nowait root /usr/lbin/ftpd    ftpd -l telnet stream tcp nowait root /usr/lbin/telnetd telnetd
    tftp   dgram  udp wait   root /usr/lbin/tftpd   tftpd bootps dgram  udp wait   root /usr/lbin/bootpd  bootpd
    finger stream tcp nowait bin  /usr/lbin/fingerd fingerd login  stream tcp nowait root /usr/lbin/rlogind rlogind
    shell  stream tcp nowait root /usr/lbin/remshd  remshd exec   stream tcp nowait root /usr/lbin/rexecd  rexecd
   ```

   To disable any of these services, comment out the line by typing a pound sign (#) as the first character on the line.

2. After modifying the `/etc/inetd.conf` file, type the following command to force `inetd` to read its configuration file:

   ```
   /usr/sbin/inetd -c
   ```

3. Make sure `/etc/inetd.conf` is owned by user `root` and group `other`, and make sure its permissions are set to 0444 (`-r--r--r--`).

For more information, type `man 4 inetd.conf` or `man 1M inetd` at the HP-UX prompt.

Editing the /var/adm/inetd.sec File

The `/var/adm/inetd.sec` file is a security file that `inetd` reads to determine which remote hosts are allowed to access the services on your host. The `inetd.sec` file is optional; you do not need this file to run the Internet Services.

To edit the `inetd.sec` file using a text editor or HP SMH, complete the following steps:

1. If the `/var/adm/inetd.sec` file does not exist on your host, copy `/usr/newconfig/var/adm/inetd.sec` to `/var/adm/inetd.sec`.
2. Create one line in `inetd.sec` for each service to which you want to restrict access. Do not create more than one line for any service.

   Each line in the `/var/adm/inetd.sec` file has the following syntax:

   ```
   service_name {allow} host_specifier [host_specifier...] {deny}
   ```

   where *`service_name`* is the first field in an entry in the `/etc/inetd.conf` file, and `host_specifier` is a host name, IP address, IP address range, or the wildcard character (*).
3. Make sure the `/var/adm/inetd.sec` file is owned by user `root` and group `other`, and make sure its permissions are set to 0444 (`-r--r--r--`).

Following are some example lines from an `inetd.sec` file:

```
login allow 10.* shell deny vandal hun tftp deny *
```

The first example allows access to `rlogin` from any IP address beginning with 10. The second example denies access to `remsh` and `rcp` from hosts `vandal` and `hun`. The third example denies everyone access to `tftp`.

Only the services configured in `/etc/inetd.conf` can be configured in `/var/adm/inetd.sec`.

For more information, type `man 4 inetd.sec` or `man 1M inetd` at the HP-UX prompt.

Configuring Logging

This section discusses the following topics:

- "Configuring syslogd" (page 33)
- "Maintaining System Log Files" (page 34)
- "Configuring inetd Connection Logging" (page 35)
- "Configuring ftpd Session Logging" (page 35)

Configuring syslogd

The Internet daemons and servers log informational and error messages through `syslog`. You can monitor these messages by running `syslogd` and determine the type and extent of monitoring through `syslogd`'s configuration file, `/etc/syslog.conf`.

Each line in /etc/syslog.conf has a selector and an action. The selector specifies which part of the system generated the message and what priority the message has. The action specifies where the message should be sent.

The part of the selector that specifies where a message comes from is called the facility. All Internet daemons and servers, except sendmail, log messages to the daemon facility. sendmail logs messages to the mail facility. syslogd logs messages to the syslog facility. You can indicate all facilities in the configuration file with an asterisk (*).

The part of the selector that specifies what priority a message has is called the level. Selector levels are debug, information, notice, warning, error, alert, emergency, and critical. A message must be at or above the level you specify in order to be logged.

The action allows you to specify where messages should be directed. You can have the messages directed to files, users, the console, or to a syslogd running on another host.

The following is the default configuration for /etc/syslog.conf:

```
mail.debug                /var/adm/syslog/mail.log *.info,mail.none         /var/adm/syslog/syslog.log
*.alert *.alert           /det/console root *.emerg                   *
```

With this configuration, all mail log messages at the debug level or higher are sent to /var/adm/syslog/mail.log. Log messages from any facility at the information level or higher (but no mail messages) are sent to /var/adm/syslog/syslog.log. Log messages from any facility at the alert level or higher are sent to the console and to any terminal where the superuser is logged in. All messages at the emergency level or higher are sent to all users on the system.

For more information about syslogd and its configuration file, type man 3C syslog or man 1M syslogd at the HP-UX prompt.

Maintaining System Log Files

The log files specified in your syslogd configuration can fill up your disk if you do not monitor their size. To control the size of these files, do the following:

1. Remove or rename your log files as in the following example:

   ```
   cd /var/adm/syslog mv mail.log mail.log.old mv syslog.log sylog.log.old
   ```

2. Restart syslogd with the following commands:

   ```
   cd /sbin/init.d syslogd stop syslogd start
   ```

When you reboot your system, each log file is moved to *filename*.old automatically, and new log files are started.

Configuring inetd Connection Logging

The inetd daemon logs connection requests through syslogd. It logs successful connections at the information level and unsuccessful connection attempts at the notice level. By default, inetd starts up with connection logging turned off.

If inetd is running with connection logging turned off, issue the following command to invoke it:

```
/usr/sbin/inetd -l
```

If inetd is running with connection logging turned on, the same command turns it off. For more information, type man 1M inetd at the HP-UX prompt.

Configuring ftpd Session Logging

To configure ftpd to log messages about an ftp session, including commands, logins, login failures, and anonymous ftp activity, complete the following steps:

1. Add the -l option to the ftp line in the /etc/inetd.conf file, as in the following example:

   ```
   ftp stream tcp nowait root /usr/lbin/ftpd ftpd -l
   ```

2. Issue the following command to force inetd to read its configuration file:

   ```
   /usr/sbin/inetd -c
   ```

For more information, type man 1M ftpd at the HP-UX prompt. This manpage contains a complete list of error messages.

See “Configuring Logging for ftp” in the *HP-UX Remote Access Services Administrator’s Guide* at the URL http://www.docs.hp.com/hpux/netcom/index.html#Internet%20Services for more information on logging ftp file transfer information.

3 TCP Wrappers

The Transmission Control Protocol (TCP) Wrappers product suite provides an enhanced security mechanism for services spawned by the Internet Services daemon, `inetd`.

This chapter discusses the following topics:

Overview

The Internet services server, `inetd`, allows a single process to wait for multiple services instead of the single process waiting for each service. When a connection is established with `inetd` for a service, `inetd` runs the appropriate server specified in the `/etc/inetd.conf` file and waits for other connections.

If you enable TCP wrappers, `inetd` runs a TCP wrapper daemon, `tcpd`, instead of running the requested service directly. When a request for a service is received, `inetd` invokes `tcpd` for the service. `tcpd` logs the request and checks the access control files for a matching daemon-client pair entry to either grant or deny access to the requested service. If access is granted to the requested service, `tcpd` invokes the appropriate server program.

You can define configuration parameters such as logging behavior, user name lookups, and reverse look up failure behavior in the configuration file `/etc/tcpd.conf.tcpd` reads the configuration file `/etc/tcpd.conf` file for configuration parameters during runtime.

The wrappers program does not work with RPC services over TCP. These services are registered as `rpc` or `tcp` in the `/etc/inetd.conf` file. The only non-trivial service that is affected by this limitation is `rexd` used by the `on` command.

The tcpd Features

The `tcpd` program provides the following features to enforce access control checks for a service:

- Access Control
- Host name or Address Spoofing
- Client User Name
- Setting Traps
- Banner Messages

Access Control

TCP wrappers uses the files /etc/hosts.allow and /etc/hosts.deny as Access Control Lists (ACLs). These access control files are used to match the client and server entries with the service request. These files are based on pattern matching and can be extended via optional extensions such as allowing spawning of a shell command.

Each access control file consists of a set of access control rules for different services that use tcpd.

An access control rule is of the following form:

```
daemon_list:client_list:option:option:...
```

daemon_list Specifies the list of daemons.

client_list Specifies the list of clients for which the access control rule is applicable. Each list is a set of items separated by a space. A client in the client_list specifies the name or address of a host requesting a service.

option Specifies a list of options. Options are separated by a colon.

The access control files are /etc/hosts.allow and /etc/hosts.deny. If you do not create these files, and specify the daemon-client pair for granting or denying access, the access control is disabled. The access control module reads these files in the following order, before granting or denying access to any service:

1. The /etc/hosts.allow file – If a daemon-client pair matches an entry in this file, access is granted.
2. The /etc/hosts.deny file – If a daemon-client pair matches an entry in this file, access is denied.
3. If a daemon-client pair match is not found in either of the access control files, access is granted.

Following are examples of different entries in the files /etc/hosts.allow and /etc/hosts.deny:

1. To grant access to the ftp service to all the users, specify the following entry in the /etc/hosts.allow file:

   ```
   ftpd:ALL
   ```

2. To deny access to the host blue.rainbow.com and all hosts in the domain rainbow.com to all the services, specify the following entry in the /etc/hosts.deny file:

   ```
   ALL:blue.rainbow.com, .rainbow.com
   ```

3. To grant the telnet service to all the hosts in the domain xyz.com except the host abc.xyz.com, specify the following entry in the /etc.hosts.allow file:

   ```
   telnetd:.xyz.com EXCEPT abc.xyz.com
   ```

For more information on the access control language and ACL options, type man 5 hosts_access or man 5 hosts_options at the HP-UX prompt.

Host Name/Address Spoofing

tcpd prevents an illegal host that behaves as a legal host from accessing services. If any discrepancy is identified in the client address or name, the wrapper program denies access to that host and logs the information. tcpd also disables the source-routing socket options on all the host's connections. This protection mechanism benefits UDP services.

Client User Name Lookup

tcpd determines the identity of a client requesting a particular TCP connection using the RFC 931 (Authentication Server) protocol. By default, the client user name lookup is disabled in the /etc/tcpd.conf configuration file. If you enable client user name lookup in the configuration file, tcpd assumes that the client requesting the service runs a RFC931-compliant daemon, such as IDENT.

Trap Setting

This feature allows you to trigger appropriate action on the host depending on the number of denied connection attempts. For example, the following rule in the /etc/hosts.deny file denies access to all hosts, and notifies when a remote host attempts to access the TFTP server:

```
tftpd:ALL:spawn (/usr/bin/safe_finger -1 @%h2>&1 | mailx -s “remote tftp attempt” root)
```

Banner Message

This feature provides a mechanism to send a message when an ACL rule is included in an access control file. For example, the following rule in the /etc/hosts.deny file sends the message in the telnetd file placed in the /tmp/banner directory, and denies access to a request from any host whose address starts with 192.5.2:

```
telnetd:192.5.2.:banners/tmp/banner
```

TCP Wrappers Files

The TCP Wrappers product suite contains the following files:

- The tcpd Daemon
- The libwrap.a Library API

- The tcpdchk Tool
- The tcpdmatch Tool
- The try-from Utility
- The safe_finger Program

The tcpd Daemon

The tcpd daemon monitors access to a service, logs the host name and the remote user name owning the connection, and performs some additional access control checks. After tcpd checks the connection, the wrapper invokes the desired server program and exits.

Enabling tcpd

You can use either of the following methods to enable tcpd:

1. Edit each entry in the /etc/inetd.conf file to include the tcpd server program, /usr/lbin/tcpd. The server program field in the /etc/inetd.conf file indicates the absolute path name of the server that inetd executes. For example, consider the following entry for the telnet service in the /etc/inetd.conf file:

   ```
   telnet stream tcp nowait root /usr/lbin/telnetd telnetd
   ```

   Edit this entry to include the tcpd server, /usr/lbin/tcpd instead of the telnet server, /usr/lbin/telnetd, as follows:

   ```
   telnet stream tcp nowait root /usr/lbin/tcpd /usr/lbin/telnetd telnetd
   ```

   NOTE: If you specify this entry without the absolute path of telnetd (/usr/lbin/telnetd), tcpd searches the telnetd binary in the /usr/lbin/wrapper directory.

   The last component of the path name, /usr/lbin/telnetd telnetd, is used for access control and logging. When the telnet service is requested, inetd invokes the tcpd server instead of invoking the telnet server. tcpd performs access control checks and verifies whether the connection is valid. Then, it invokes the telnetd server. Similarly, you can change the entries for other services to include the tcpd server.

   After making the relevant changes, you must reconfigure inetd using the following command on the command line:

   ```
   # inetd -c
   ```

2. Move the service daemon to the /usr/lbin/wrapper directory, and move tcpd to the location of the service daemon. You need not make any changes to the

/etc/inetd.conf file. For example, you can enable the ftpd service with tcpd by executing the following commands at the command prompt:

```
# mkdir /usr/lbin/wrapper

# mv /usr/lbin/ftpd /usr/lbin/wrapper

# cp tcpd /usr/lbin/ftpd
```

When an ftp service is requested, inetd spawns the /usr/lbin/ftpd daemon which is actually the tcpd daemon. Then, tcpd performs access control checks before invoking the ftpd daemon in the /usr/lbin/wrapper directory.

For more information on tcpd configuration, type man 1M tcpd or man 4 tcpd.conf at the HP-UX prompt.

The libwrap.a Library

The libwrap.a library provides a set of APIs for independent applications to enforce host access control based on the files /etc/hosts.allow and /etc/hosts.deny files. The APIs implement a rule-based access control language with optional shell commands, that are executed when a rule is invoked.

To enforce the host access control in an independent daemon, a service must include the tcpd.h header file and link with the libwrap.a library APIs. The libwrap.a library contains the following APIs:

- request_init()

  Initializes the request_info structure with the client request information.

- request_set()

  Updates an initialized request_info structure.

Both the APIs request_init() and request_set() accept the request_info structure and a variable length list of key-value pairs as input parameters; and, they return the request_info structure defined in the tcpd.h header file. The argument lists are terminated with a zero key value.

- hosts_access()

  Reads the ACLs and returns either 1 or 0 indicating the access granted or denied, respectively.

- hosts_ctl()

  This API is a wrapper to the functions request_init() and hosts_access(). It accepts the daemon name, client host name, client address, and user name as input parameters. The client host name, address, and user name arguments must contain valid data or STRING_UNKNOWN defined in the `tcpd.h` file. If the access is denied the `hosts_ctl()` API returns a value `0`.

The following are the methods to implement access control checks in a daemon program:

1. Fill the variable elements in the structure `request_info` using the routines `request_init()` and `request_set()`, and call the `hosts_access()` routine to verify these elements with the ACLs.
2. Call the function `hosts_ctl()` with appropriate input parameters to check with the ACLs.

For more information on these APIs, type `man 3 hosts_access` at the HP-UX prompt.

The tcpdchk Tool

The `tcpdchk` tool performs the following functions:

- Examines the validity of entries in the `/etc/inetd.conf` file and ACLs.
- Inspects the TCP wrapper configurations and reports problems, if any.
- Checks the `tcpd` access control files (`/etc/hosts.allow` and `/etc/hosts.deny`), and compares the entries in these files with the entries in the `/etc/inetd.conf` file.

You can run the `tcpdchk` tool on the command line as follows:

```
/usr/bin/tcpdchk [-a] [-d] [-i inet_conf] [-v]
```

where,

`-a` Reports access control rules that grant access without an explicit `ALLOW` keyword.

`-d` Examines the files `/etc/hosts.allow` and `/etc/hosts.deny` in the current directory instead of the default directory.

`-i inet_conf` Specifies a different location for the configuration file `inetd.conf` instead of the default directory, `/etc/inetd.conf`.

`-v` Displays the contents of an access control rule, that is, the daemon list, client list, shell command and option, in a printable format. This helps you to identify discrepancies in the output.

For more information, type `man 1 tcpdchk` at the HP-UX prompt.

The tcpdmatch Tool

The `tcpdmatch` tool simulates the wrappers daemon program, and `tcpd` behavior for a particular host and service.

`tcpdmatch` predicts how the TCP wrapper daemon handles a specific service request. It examines the `tcpd` access control tables (`/etc/hosts.allow` and `/etc/hosts.deny`) and prints a report. For maximum accuracy, it extracts additional information from the `/etc/inetd.conf` file.

You can execute the tcpdmatch tool on the command line using the following formats:

1. `/usr/bin/tcpdmatch [-d] [-i inet_conf] daemon client`

2. `/usr/bin/tcpdmatch [-d] [-i inet_conf] daemon@[server] [user@]client`

daemon Specifies a daemon name.

client Specifies the host name, network address, or the unknown or paranoid wildcard formats.

server Specifies a host name or network address or the unknown or paranoid wildcard formats.

user Indicates a client user identifier, and specifies a login name or address. The default user name is unknown.

You can use the first tcpd syntax when a server has more than one address or name.

Example 3-1 Sample Usage of the tcpdmatch Tool

The following example denotes how tcpd handles an ftp request from a local system:

```
# tcpdmatch ftpd localhost
```

If the host name lookup fails, the same request is handled by tcpd as follows:

```
# tcpdmatch ftpd 127.0.0.1
```

To determine what tcpd does when the client name and address do not match, execute the following command at the command prompt:

```
# tcpdmatch ftpd paranoid
```

For more information on the tcpdmatch tool, type man 1 tcpdmatch, and for more information on wildcard patterns, type man 5 hosts_access at the HP-UX prompt.

The try-from Utility

The try-from utility identifies the detailed information regarding to a connection. This utility must be called via a remote shell command to determine if the host name and the address are recognized properly, and also if the user name lookup works.

The try-from utility can be executed from the command line as follows:

```
# remsh host /usr/bin/try-from
```

When the try-from utility is invoked, it prints the following output:

```
client address       (%a): client hostname      (%n): client username      (%u): client info          (%c):
server address       (%A): server hostname      (%N): server process       (%d): server info            (%s):
```

The client information describes how the remote host recognizes the client in terms of an address, name, and user name, whereas, the server information describes the remote host.

For more information on % <letter> expressions, type man 5 hosts_access at the HP-UX prompt.

The safe_finger Program

safe_finger, a wrapper program to the finger client, protects the data sent by the remote finger server. This program accepts all the options supported by the finger client.

For more information, type man 1 finger at the HP-UX prompt.

Following is an example of the safe_finger command:

```
# /usr/bin/safe_finger -l @xyz.abc.def.com
```

This command prints the user information on the remote host xyz.abc.def.com.

HP recommends you to use this program to implement traps in the access control language of the files /etc/hosts.allow and /etc/hosts.deny.

For more information on setting traps, type man 5 hosts_access at the HP-UX prompt.

IPv6 Support

To enable access control mechanism to IPv6 connections of a service, you must enable IPv6 support to that service in the /etc/inetd.conffile. You must specify the protocol in the /etc/inetd.conf file as tcp6 or upd6, to enable IPv6 support for a particular service.

For example, to support IPv6 functionality for the ftpd service, you must modify the /etc/inetd.conf file as follows:

```
ftp stream tcp6 nowait root /usr/bin/ftpd ftpd -l
```

However, if you specify tcp instead of tcp6, ftp operates in the IPv4 mode.

Troubleshooting

tcpd logs the connection-related information and problems encountered during a connection in the /var/adm/syslog/syslog.log file, before invoking the actual service daemon.

You can enable logging in tcpd by specifying the logging level parameter in the /etc/tcpd.conf file. The syntax for specifying the logging level parameter is as follows:

```
log_level { normal | extended }
```

With an `extended` parameter, `tcpd` logs the ACL information, such as, the entry with which the client request is matched and the entry's related options. The default logging level parameter is `normal` which logs connection details, such as, acceptance or refusal of connections.

TCP wrappers provides the tools `tcpdchk` and `tcpdmatch` for troubleshooting. `tcpdchk` validates the `inetd.conf`, `hosts.allow` and `hosts.deny` entries in the configuration file, and `tcpdmatch` determines how `tcpd` handles a specific service request.

For more information on the `tcpdchk` and `tcpdmatch` tools, see "The tcpdchk Tool" (page 42) and "The tcpdmatch Tool" (page 42), respectively.

4 Configuring NTP

The Network Time Protocol (NTP) assures accurate synchronization of the computer's clock time with reference to a number of primary reference sources, using an equipment such as a radio receiver. NTP runs as a continuous background client process on a system, and sends periodic time requests to primary servers to obtain the time stamps. It also checks for errors caused due to equipment or propagation failures.

This chapter describes the basic and advanced NTP concepts, components and configuration instructions required to use NTP. This chapter also includes troubleshooting information.

This chapter contains the following topics:

The section Getting Started with NTP is ideal for beginners with limited experience on NTP. The section Advanced NTP Topics is ideal for anexperienced user with sufficient experience on NTP.

Getting Started with NTP

The Network Time Protocol (NTP) is a family of programs used to adjust the system clock on your computer and to synchronize it with external sources of time. Computers are very sensitive to time deviations caused by drifting. All clocks drift including the clock inside the computers. NTP provides accurate time in the range of microsecond to millisecond and helps overcome drifting.

Some of the pervasive computing processes that may be affected by disparity in time include: debugging, database and transaction processing, and compiling software using the `make` utility.

Debugging system problems becomes difficult if the timestamp in the system logs are not true.

Databases rely on time to a large extent. Databases and transaction processing application may get confused if clients and servers have different times.

The `make` utility is commonly used to manage the compilation of software. This utility verifies file timestamps, with one-second granularity, to decide which `.O` files need to be rebuilt when the underlying source files are changed. The problem increases when files on machines, at various sites in different time zones, need to be compiled and built into the new version of the source file. Also, if some directories are NFS mounted, and the server and client have different notions of the current time, `make` can fail to rebuild some derived objects. This can produce an executable that is not based on the latest sources. A one-second granularity of file stamp indicates that the client and server must be synchronized close to 1000 milliseconds, to ensure that `make` compiles the appropriate files.

The following topics are discussed in this section:

- "NTP Equipment" (page 48)
- "Choosing the Source of Time" (page 48)
- "Backup Time Servers" (page 56)
- "Configuring Your Primary NTP Server" (page 58)

NTP Equipment

The following equipments are required to effectively use the NTP programs:

- Internet or your own radio receiver, such as GPS (Global Positioning System), as a time source.
- An ordinary network, such as an Ethernet, in your building.
- Familiarity with configuring and setting up NTP.

Starting NTP Configuration

For a basic NTP configuration, you must complete the following steps:

1. Choose a source of time.
2. Determine how frequently the system clock must synchronize with the source of time.
3. Select backup time server.
4. Configure the primary NTP server.

The following sections cover these steps in detail.

Choosing the Source of Time

Government organizations define, regulate and distribute time for synchronization purpose. These organizations constantly coordinate and synchronize their clocks with other organizations within nanoseconds of each other. The first step in using NTP is selecting the best source of time for your organization.

You must be careful while selecting the source of time. If the type of applications and processes the users run are sensitive to time, you must select the source of time that provides a stable time, and is not affected by network delays. Do not select the NTP service depending on price.

Also, select a source of time that you can reach fast. The closer the source of time, the better. Choose a source that is physically close and one that takes few network hops to reach the destination. For more information on physical and network distance, see "Configuring Mulltiple Time Servers" on page 221.

Available Time Sources

The most common time distribution mechanisms from which you can draw time are:

- Public time server (phone or modem) via the Internet
- Local clock impersonators
- Radio receiver – Terrestrial and satellite broadcast

Public Time Server

You can connect to public time servers via the Internet free of charge for a limited time. Public time servers also provide dial-up access through a modem. This is the cheapest and most popular method. One of the main disadvantage of this option is that all the networks cannot use the public time server because they are protected behind firewalls.

Many other time servers are available that you can access. HP provides a public time server, which is located in Cupertino, California. You can use this time server if you are located in North America. The following lists the details for this time server:

```
ntp-cup.external.hp.com (192.6.38.127) Location: Cupertino, CA (SF Bay Area) 37:20N/122:00W
Synchronization: NTP3 primary (GPS), HP-UX Service Area: West Coast USA Access Policy: open access
Contact: timer@cup.hp.com Note: no need to notify for access, go right ahead!
```

NOTE: An enterprise can implement its own hierarchy of NTP time servers, including the stratum-1 servers. If your administrative domain is part of an enterprise-wide Internet, you must verify for available NTP resources in your enterprise. If your administrative domain does *not* have access to lower-stratum time servers, NTP servers on the Internet are available that provide public time synchronization. You can use stratum-1 and stratum-2 servers only with the permission of the administrator; you must always check with the administrator before using an NTP server on the Internet.

Local Clock Impersonators

You can use a local clock impersonator in either of the following instances:

- If you are behind a firewall.
- If you are not connected to the Internet.
- If you cannot afford a radio receiver.

You can declare your NTP machine as a time server, and this machine can serve time within a closed domain. Because this time server is isolated, it does not synchronize with the real time.

IMPORTANT: Using this option may cause problems if you are always connected outside your domain.

To set up the local clock impersonator, add the following entry to the `/etc/ntp.conf` file:

```
server 127.127.1.1   minpoll 3   maxpoll 4
```

Radio Receiver

The radio receiver is the most accurate and expensive time distribution mechanism. Radio receiver provides a stable time and is not affected by network delays, congestion, or outrages. Some of the popular radio receiver methods are: GPS (Global Positioning System), WWV (Terrestrial North America), and DCF77 (Terrestrial Europe).

You must consider the cabling options when you select the radio receiver. Antenna cables are very expensive and RS232 cabling has a limited range.

The official HP supported GPS receivers are HP58503 driver#26 and Trimble Palisade driver#29. The only supported WWVB receiver is Spectracom Netclock/2 driver#4. DCF77 (AM and FM) signals radiate from Frankfurt Germany. DCF77 receivers are not officially supported by HP.

Setting up an HP58503A GPS Receiver

The following steps describe how to set up an HP58503A GPS:

1. Install and connect the receiver and antenna to a serial port on the HP-UX machine.
2. Append the following entries in the `/etc/ntp.conf` file:

   ```
   server  127.127.26.1  minpoll  3  maxpoll  4 # fudge 127.127.26.1  time1 -0.955  #s700
   # fudge 127.127.26.1  time1  -0.930  #s800
   ```

3. Uncomment the appropriate `# fudge` entry for your architecture. Uncomment the `#fudge ... #s800` entry for servers or uncomment `#fudge ... #s700` entry for workstations.

Setting up a Trimble Palisade GPS Receiver

The following steps describe how to set up a Trimble Palisade GPS receiver:

1. Install and connect the receiver and antenna to a serial port on the HP-UX machine.
2. Append following entries in the `/etc/ntp.conf` file:

   ```
   server  127.127.29.1  #poll period is fixed at 32 seconds # no fudge required
   # fudge 127.127.26.1  time1  -0.930  #s800
   ```

3. Create the following symbolic link:

   ```
   /usr/bin/ln -s /dev/tty0p0  /dev/palisade1
   ```

Setting up a Spectracom Netclock/2

The following steps describe how to set up a Spectracom Netclock/2:

1. Install and connect the WWVB receiver to a serial port on the HP-UX machine.
2. Append the following entries in the `/etc/ntp.conf` file:

```
server  127.127.4.1  minpoll  3  maxpoll  4 # no fudge required # fudge 127.127.26.1  time1  -0.930  #s800
```

3. Create the following symbolic link:

```
/usr/bin/ln -s /dev/tty0p0  /dev/wwvb1
```

Location of Time Source

You must always select a time server that is physically close to your network; otherwise, it may lead to poor network connectivity and delays. You must also consider the network path that a packet needs to travel, because if a time server is physically close but takes excessive number of hops to reach, you may experience network delays.

If applications on the network need to be accurate to the millisecond, you must consider the dispersion measurements and the network service quality. *Dispersion* is a measurement of the time server quality and network quality.

NOTE: Dispersion is high if the network is slow or overloaded, irrespective of the quality of the time server of network.

The time server that returns a quick response to the `ping` command is the most appropriate time server.

Figure 4-1 shows the time difference between the NTP client and the NTP time servers situated physically in different places.

Figure 4-1 Survey of Best Time Servers

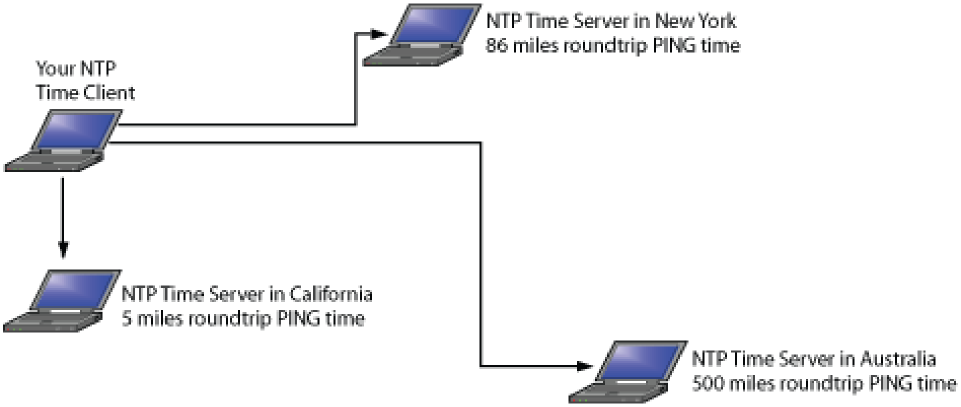
NTP Time Server in New York
86 miles roundtrip PING time
Your NTP
Time Client
NTP Time Server in California
5 miles roundtrip PING time
NTP Time Server in Australia
500 miles roundtrip PING time

The best primary server for the NTP client located in California is the time server situated in New York because the ping command response time is only 5 millisecond. The `ping` command response time for the time server in Australia takes 500 milliseconds. Therefore, selecting the time server situated in Australia is not recommended because it may cause network delays.

Example 1: Locating the Best Primary Server

Table 4-1 shows the servers the time client can access. The primary time server is NAVOBS1.MIT.EDU. The other time servers within reasonable physical and network distance are cs.columbia.edu, 129.236.2.199, and clepsydra.dec.c .

Table 4-1 Available Time Servers

```
remote            refid          st t  when poll  reach delay   offset   disp
==============================================================================
clepsydra.dec.c   usno.pa-x.dec  2  u  927  1024  355   108.49  -18.215   3.63
*NAVOBS1.MIT.EDU  .USNO.         1  u  214  1024  377   38.48   -0.536    0.90 ticks.CS.UNLV.ED  to
ck.CS.UNLV   3  u  721  1024  377   2113.97  1004.94  824.57 -cunixd-ether.cc  192.5.4
1.209   2  u  636  1024  377   47.99    3.090    9.75
+cs.columbia.edu  haven.umd.edu  2  u  172  1024  377   3.39     12.573   1.14
+129.236.2.199    BITSY.MIT.E    2  u  423  1024  376   13.43   -14.707   22.60
```

Choose three (or more) time servers that are geographically close to the NTP client. If you are located in London, it is not preferable to choose time servers in Australia or Brazil. Long distances over water usually pose a poor network connection due to delay and path symmetry. Router hops also delay the packets in unpredictable ways.

Before using a time server, you must evaluate these potential time servers (and the network paths) to decide if they are physically close to the NTP client (consider ping time, delay and variation) and configured properly. You may also have to send notification to the time server before using them (consider the ettiquitte of the listings at UDelaware). Do not point more than three machines to a single public time server. Use a small group of machines (at stratum-2 or stratum-3) as the main time server for remaining systems. For more information about stratum levels, see “Stratum Levels and Time Server Hierarchy” (page 59).

The public stratum-2 servers provides time service for all clients. Also, their access policies are less restrictive than the stratum-1 servers. The errors displayed while connecting your machine with the public time server (or ISP) denotes the quality of the network service.

This makes the distinction between stratum-1 and stratum-2 almost meaningless for most purposes.

Unlike other applications (such as FTP, DNS, NFS, Sendmail) which can tolerate huge delays in packet, NTP is sensitive to the network service quality. Therefore, in addition to the quality of the time servers, you must ensure that the network is fast and is not overloaded. A minor delay in the network can be harmful for your time service. Delays shows up immediately in the dispersion figures.

If you can afford time delays in the time service upto milliseconds, then pay special attention to the network service quality. If you can afford time delays in time service upto microseconds, you must abandon the network time servers and choose a radio clock for each NTP client.

You can evaluate different public time servers from the stratum-2 list.

Following is the stratum-2 listing of the an HP machine which was provided in the Silicon Valley for public use in North America.

```
ntp-cup.external.hp.com (192.6.38.127) Location: Cupertino CA (SF Bay area) 37:20N/122:00W
Synchronization: NTPv3 primary (GPS), HP-UX Service Area: West Coast USA Access Policy: open access
Contact: timer@cup.hp.com Note: no need to notify for access, go right ahead!
```

If you are located in Silicon Valley, you can `ping` this time server and notice that the time server is 5 milliseconds.

```
/usr/sbin/ping ntp-cup.external.hp.com 64 5     PING ntp-cup.external.hp.com: 64 byte packets
    64 bytes from 192.6.38.127: icmp_seq=0. time=5. ms     64 bytes from 192.6.38.127: icmp_seq=1. time=4. ms
    64 bytes from 192.6.38.127: icmp_seq=2. time=4. ms     64 bytes from 192.6.38.127: icmp_seq=3. time=5. ms
    64 bytes from 192.6.38.127: icmp_seq=4. time=5. ms ----ntp-cup.external.hp.com PING Statistics---
   5 packets transmitted, 5 packets received, 0% packet loss    round-trip (ms) min/avg/max = 4/4/5
```

Determining Synchronization Sources

You can query the time server using the following command to check the synchronization sources:

```
/usr/bin/ntpq -p ntp-cup.external.hp.com
```

Table 4-2 displays the synchronized time servers.

Table 4-2 Locating Synchronized Time Servers

```
remote           refid      st t  when   poll  reach  delay   offset    disp
==============================================================================
*REFCLK(29,1)     .GPS.       0  l   35      32    376    0.00   -0.004    0.02
-bigben.cac.wash .USNO.      1  u   47     128    377   40.16   -1.244    1.37
clepsydra.dec.c  usno.pa     2  u   561   1024    377   16.74   -4.563    4.21
-clock.isc.org   .GOES.      1  u   418   1024    377    6.87   -3.766    3.57
hpsdlo.sdd.hp.c  wwvb.col. 2  u   34     16      204   48.17   -8.584  926.35
+tick.ucla.edu   .USNO.      1  u   111    128    377   20.03   -0.178    0.43
+usno.pa-x.dec.c .USNO.        1   u    42    128 377    6.96   -0.408    0.38
```

This time server is synchronized (asterisk in column one) to `REFCLK(29,1),` which is a Trimble Palisade GPS receiver. The offset from GPS is 0.004 milliseconds and the dispersion is 0.02 milliseconds. This time server also has several good stratum-1 and stratum-2 servers which it can fall back if the GPS receiver stops functioning.

The entry for the time server `hpsdlo.sdd.hp.com` has a delay, offset, and dispersion measures less than any of the other sources. The time server `hpsdlo` is good, but the network in between has some problems, depicted by the large dispersion figures. If the dispersion numbers are high, contact the network service provider.

The time server `ntp-cup.external.hp.com` is the appropriate time server because it is only 5 milliseconds from the NTP client and it the right choice for a public time server. The `ping` command round-trip time determines whether this time server is the right choice for a public time server.

Example 2: Evaluating Time Servers in Eastern United States

For a time server located on the east coast of Northern America, following are the details:

```
ntp.ctr.columbia.edu (128.59.64.60) Location: Columbia University Center for Telecommunications Research; NYC
Synchronization: NTP secondary (stratum 2), Sun/Unix Service Area: Sprintlink/NYSERnet
Access Policy: open access, authenticated NTP (DES/MD5) available Contact: Seth Robertson (timekeeper@ctr.columbia.edu)
 Note: IP addresses are subject to change; please use DNS   /usr/sbin/ping ntp.ctr.columbia.edu 64 5
   PING 128.59.64.60: 64 byte packets           64 bytes from 128.59.64.60: icmp_seq=0. time=83. ms
           64 bytes from 128.59.64.60: icmp_seq=1. time=86. ms
           64 bytes from 128.59.64.60: icmp_seq=2. time=85. ms
           64 bytes from 128.59.64.60: icmp_seq=3. time=86. ms
           64 bytes from 128.59.64.60: icmp_seq=4. time=83. ms          ----128.59.64.60 PING Statistics----
   5 packets transmitted, 5 packets received, 0% packet loss          round-trip (ms)  min/avg/max = 83/84/86
```

In this example, the ping round-trip times are significantly greater than in the previous example. 85 milliseconds is good for general NTP purposes. The dispersion measurements are less than the `ping` round-trip times. The NTP daemon has a watershed at 128 milliseconds, but this example server at 85 milliseconds is comfortably below that. You can use the server at Columbia.

```
/usr/sbin/ntpq -p ntp.ctr.columbia.edu
```

Table 4-3 describes time servers in eastern United States.

Table 4-3 Evaluating Time Servers in Eastern United States

```
   remote           refid      st t when poll reach   delay   offset    disp
==========================================================================
+clepsydra.dec.c usno.pa-x.dec.c  2 u  927 1024  355   108.49  -18.215    3.63
otc1.psu.edu    .WWV.           1 -  17d 1024    0    28.26  -25.362 16000.0
*NAVOBS1.MIT.EDU .USNO.          1 u  214 1024  377    38.48   -0.536    0.90
tick.CS.UNLV.ED tock.CS.UNLV.ED  3 u  721 1024  377  2113.97  1004.94  824.57
132.202.190.65  0.0.0.0         16 -    - 1024    0     0.00    0.000 16000.0
unix.tamu.edu   orac.brc.tamus.  3 u  636 1024  377    47.99    3.090    9.75
at-gw2-bin.appl 0.0.0.0         16 -    - 1024    0     0.00    0.000 16000.0
-cunixd-ether.cc 192.5.41.209     2 u  172 1024  377     3.39   12.573    1.14
cunixd.cc.colum 0.0.0.0         16 u  285   64    0     0.00    0.000 16000.0
+cs.columbia.edu haven.umd.edu    2 u  906 1024  376     2.41   -5.552   15.12
+129.236.2.199   BITSY.MIT.EDU    2 u  423 1024  376    13.43  -14.707   22.60
 cucise.cis.colu cs.columbia.edu 3 u   62 1024  377     5.84   -1.975   12.70
```

This time server at Columbia University has a variety of stratum-1, stratum-2, and stratum-3 sources, which is good. It also has three sources which are not responding right now (reach=0), and one with very large delay, offset, and dispersion (tick.CS.UNLV.EDU). As before, this is due to networking problems between client and server (New York to Las Vegas, over 3000 km), not some fault with the NTP implementation at either end. This time server at Columbia is currently synchronized to NAVOBS1.MIT.EDU, but three others (marked with "+" in column one) are attractive and could step in immediately if NAVOBS1 failed for any reason.

Example 3: Evaluating Time Servers in Australia

Look at a time server in Australia. Here are the details:

```
ntp.adelaide.edu.au (129.127.40.3)  Location: University of Adelaide, South Australia
Synchronization: NTP V3 secondary (stratum 2), DECsystem 5000/25 Unix  Service Area: AARNet  Access Policy: open access
 Contact: Danielle Hopkins (dani@itd.adelaide.edu.au)
```

```
/usr/sbin/ping  ntp.adelaide.edu.au 64 5
```

```
PING huon.itd.adelaide.edu.AU: 64 byte packets      64 bytes from 129.127.40.3: icmp_seq=0. time=498. ms
     64 bytes from 129.127.40.3: icmp_seq=1. time=500. ms      64 bytes from 129.127.40.3: icmp_seq=2. time=497. ms
      64 bytes from 129.127.40.3: icmp_seq=3. time=498. ms      64 bytes from 129.127.40.3: icmp_seq=4. time=496. ms
```

```
----huon.itd.adelaide.edu.AU PING Statistics----
```

```
5 packets transmitted, 5 packets received, 0% packet loss round-trip (ms)  min/avg/max = 496/497/500
```

Assume you are located in western United States and you `ping` this time server. The ping round-trip times are much larger; around 500 milliseconds. Do not use a time server at this distance unless you are really need it and understand what 500 milliseconds step changes mean to your users and applications. However, depending on your location, `ping` round trip times from this server may be acceptable levels. The round-trip times from your own location might be much smaller. Also note that the variation in round-trip times is small.

```
/usr/sbin/ntpq -p ntp.adelaide.edu.au
```

Table 4-4 Evaluating Time Sources in Australia

```
     remote           refid      st t when poll reach   delay   offset    disp
==============================================================================  .otto.bf.rmit.ed 130.155.98.1     2 u  229 1024
  376    16.34    7.132    7.87  .student.ntu.edu murgon.cs.mu.OZ  2 u   47  128  377    81.34    5.166    5.25  .203.31.96.1
    murgon.cs.mu.OZ  2 u   13  256  373   115.74   30.147   38.54  .203.172.21.222  tick.usno.navy.  2 u   43 1024  367   866.64
   47.316   65.32  -128.184.1.4     tictoc.tip.CSIR  2 u   99  128  377    13.40   -2.976    5.66  129.127.40.255  0.0.0.0
     16 u    -   64    0     0.00    0.000 16000.0  *tictoc.tip.CSIR .ATOM.           1 u   17   64  377    26.92   -0.071
1.71  .dishwasher1.mpc gilja.itd.adela  3 u  164  256  376    35.78    4.769    5.66  xclepsydra.dec.c usno.pa-x.dec.c  2 u
1468 1024  376   473.36  -53.841   12.89  murgon.cs.mu.OZ .GPS.           1 u  47d 1024    0    16.19  -398.80 16000.0
-augean.eleceng. murgon.cs.mu.OZ  2 u   12  128  377     1.83    3.270    1.21  .ns.saard.net    augean.eleceng.  3 u   27
64  375    0.92   -0.013    1.19  +cuscus.cc.uq.ed tictoc.tip.CSIR  2 u   28   64  376    34.91    1.981    1.27  +staff.cs.usyd.e
 tictoc.tip.CSIR  2 u    3   64  375    25.21    0.158    1.97  .wasat.its.deaki tictoc.tip.CSIR  2 u    1  128  377    15.37
  -2.492    1.69  .luna.its.deakin tictoc.tip.CSIR  2 u  123  128  172    16.11   -0.350  501.11  -earth.its.deaki tictoc.tip.CSIR
  2 u   28  128  377    12.19   -3.582    2.15  phobos.its.deak tictoc.tip.CSIR  2 u  169  128   56    12.42   -2.325 1000.76
 .sol.ccs.deakin. tictoc.tip.CSIR  2 u  136  512  265    13.89   -1.083  251.83  +argos.eleceng.a tictoc.tip.CSIR  2 u   23
64  377    1.82    0.197    1.21  .mercury.its.dea tictoc.tip.CSIR  2 u  123  256  377    16.91   -2.584    2.94  .orion.atnf.CSIR
 murgon.cs.mu.OZ  2 u  111  512  376    53.51   -0.712    5.92  +smig2a.City.Uni tictoc.tip.CSIR  2 u   49   64  376     7.14
   0.268    1.07  +svdpw.City.UniS murgon.cs.mu.OZ  2 u   26   64  376    4.90   -0.833    1.88  .news.nsw.CSIRO. murgon.cs.mu.OZ
  2 u   54 1024  377   135.85   43.108   62.45  +210.8.40.225    murgon.cs.mu.OZ  2 u    2   64  377    50.83    1.811   14.45
  .203.103.99.66    tictoc.tip.CSIR  2 u  342 1024  376    82.82  -14.124   36.21  xpellew.ntu.edu. tictoc.tip.CSIR  2 u  408
1024  377   404.33  -159.77  161.36  xxox.lifelike.co tick.usno.navy.  2 u  494 1024  377   504.56  -59.200    5.60
```

This time server in Australia has one excellent stratum-1 source (`tictoc.tip.CSIR`) which it is currently synchronized to, one stratum-1 source which has not responded in a while (`reach=0`), and a wide selection of stratum-2 sources (desirable candidates marked with +). Some of the stratum-2 sources are less attractive due to high delay, offset, and dispersion numbers. They are marked `x` in column one.

This time server in Australia might be a good choice for you if you are reasonably nearby, but it is probably not a good choice for time clients in North America.

When the time server in Silicon Valley is configured to use `sirius.ctr.columbia.edu` and `gpo.adelaide.edu` as time sources, the output from `ntpq -p` looks like this (about 10 minutes after daemon startup):

Table 4-5 Output from ntpq for Configuring Silicon Valley Time Server

```
     remote      refid      st t when poll reach   delay   offset    disp
=========================================================================
*REFCLK(29,1)    .GPS.        0 l   25   32  377     0.00    0.413    0.03
+bigben.cac.wash .USNO.       1 u   56   64  377    39.54   -0.466    1.68
clepsydra.dec.c usno.pa-x.  2 u  122  512  377     6.32   -0.250    0.92
-clock.isc.org   .GOES.       1 u  149  512  357     5.98   -3.045    0.46
hpsdlo.sdd.hp.c wwvb.col.h  2 u   25   32  126    56.29   -8.078    8.50
+tick.ucla.edu   .USNO.       1 u   13   64  177    19.29   -0.265    0.26
+usno.pa-x.dec.c .USNO.       1 u   56   64  277     6.82    0.034    0.20
gpo.adelaide.ed tictoc.tip  2 u   15   16  377   470.52   54.789    0.90
sirius.ctr.colu NAVOBS1.MI  2 u    3   16  377    83.37   -8.372    1.24
```

The time server in Australia has a delay of 470 milliseconds, which is very similar to the `ping` round-trip times seen earlier. This leads to an offset value of 54 milliseconds, which is significantly worse than any of the other time sources. It is interesting to note that the offest is much less than the delay, which means that the round-trip is almost symmetric. NTP must assume the outbound and inbound travel times are equal, and the offset value gives an idea how unequal they might be. This is considerably better than 470/2 which would be the offset if NTP did not make this assumption. Also interesting is the very low dispersion value, which means that the round-trip time does not vary a lot as more packets are exchanged. Less than 1 millisecond is an excellent dispersion value for a trip of 15,000 kilometers. The time server in Australia is working out better than expected at this distance, but it is still noticeably poorer than the other choices that are in North America.

The time server at Columbia is better than the time server in Australia, due to the closer distance, but still noticeably worse than all of the other time sources.

You must choose a minimum of one time server, and it is a good idea to choose three or more for redundancy. Then put lines like this at the end of your `/etc/ntp.conf` file:

```
server ntp-cup.external.hp.com server bigben.cac.washington.edu server sirius.ctr.columbia.edu
```

Backup Time Servers

After selecting the primary time server, you must select two additional time servers that serve as backup time servers. The closest and fastest time server must be the primary time server. Backup time servers act as stand-by servers when the primary time server is not available. The process of establishing backup servers is known as employing redundancy. Backup time servers ensure that time sensitive applications have an alternative reliable source for time synchronization.

NOTE: You should select at least three backup servers for redundancy.

The NTP Daemon

The daemon, `xntpd`, is the network time protocol daemon. It is an operating system daemon that sets and maintains the system time in synchronization with the Internet standard time servers. `xntpd` is an implementation of NTP Version 3, as defined in

RFC 1305 (Network Time Protocol Version 3 – Specification and Implementation). It is also compatible with the NTP servers Version 1 and 2 as defined in RFC 1059 (Network Time Protocol Version 1 – Specification and Implementation) and RFC 1119 (Network Time Protocol Version 2 – Specification and Implementation), respectively.

`xntpd` operates in the following modes: symmetric active, symmetric passive, client/server, broadcast, and multicast mode, as specified in RFC 1305. Abroadcast or multicast client can deploy a group of workstations without specifying configuration details specific to the local environment A broadcast or multicast client can perform the following functions:

- Discover remote server.
- Compute client/server propagation delay correction factors.
- Configure itself automatically.

`xntpd` reads the NTP configuration file, `/etc/ntp.conf`, during startup to determine the synchronization sources and operating modes. You can also specify the configuration options on the command line when you start `xntpd`. While `xntpd` is running, you can also display `xntpd` variables and modify configuration options using the `ntpq` and `xntpdc` utilities. For more information, type `man 1M xntpd`, `man 1M ntpq` or `man 1M xntpdc` at the HP-UX prompt.

The NTP Configuration File

The NTP configuration file, `/etc/ntp.conf`, contains the initial values for synchronization sources, modes and other related information. The `xntpd` daemon reads the `/etc/ntp.conf` file during startup to determine the initial configuration values for the time server. The configuration file format is similar to a UNIX® configuration file. You can insert comments with a pound symbol (#) in the beginning of the line. A configuration command consists of an initial keyword followed by a list of arguments. Arguments can be host names, host addresses, integers, floating point numbers and text strings.

The configuration commands for a `peer`, `server` and `broadcast` are as follows:

```
peer address [ key key_id ] [version version _id ] [ prefer ]

server address [ key key_id ] [ version version_id ] [ prefer ] [ mode mode ]

broadcast address [ key key_id ] [ version version_id ] [ ttl ttl ]
```

You can use these commands to specify either the name or address of the time server, and the mode in which the time server must operate. For more information, type `man 1M xntpd` at the HP-UX prompt.

Configuring Your Primary NTP Server

The following steps describe how to configure the primary NTP server:

1. Install the latest version of NTP on the system.
2. Select a source of time: radio receivers, public time server or local NTP system.
3. Add the server name to the file `/etc/ntp.conf` using the following command:

   ```
   server my_server.mydomain.my_org.com
   ```

   `my_server.mydomain.my_org.com`is the complete name of the server.
4. Specify the time source and add its information to the configuration file.
   - For Radio Receivers:
     a. Uncomment the following `fudge` line found at the end of the file `/etc/ntp.conf` server 127.127.26.1.

        ```
        #fudge 127.127.26.1 time1 -0.955
        ```

     b. Make a link to the device file that corresponds to the serial port you are connecting to the GPS unit by typing the following: `/usr/bin/ln -s /dev/tty0p0 /dev/hpgps1`(device name for HP GPS)
   - For the local NTP Machine, add the following lines at the end of the `/etc/ntp.conf` file:

     ```
     server 127.127.1.1
     fudge 127.127.1.1 stratum 10
     ```

     Create a link to the device file that corresponds to the serial port you are connecting to the GPS unit by adding the following line to the device file: `/usr/bin/ln -s /dev/tty0p0 /dev/hpgps1`

     Use this option only when NTP is used in an isolaed environment without a radio clock, NIST modem or Internet connection. You can also use this if a particular server clock will be used as a last resort, when all other normal synchronization sources are not availiable.
5. Start the NTP daemon using the following steps:

   a. Edit the `/etc/rc.config.d/netdaemons` file. Set the variable `NTPDATE_SERVER` equal to an NTP time server that is reachable. For example:

   ```
   NTPDATE_SERVER=15.13.108.1
   ```

   This will run the `/usr/sbin/ntpdate` command just before the NTP daemon is started, and bring your system clock close to the other server to start.

   b. Set the `XNTPD` variable to `1`.

   This starts the daemon automatically when the system transitions from level 1 to 2.

c. Start the daemon using the startup script:

```
/sbin/init.d/xntpd     start
```

d. Verify whether the daemon process is running using the following command:

```
ps -ef   grep ntp
```

The line `/usr/sbin/xntpd` appears in the list of running processes.

Advanced NTP Topics

This section includes advanced NTP topics and is ideal for experienced users. The following sections are covered in this section:

- "Stratum Levels and Time Server Hierarchy" (page 59)
- "Configuring NTP using the Configuration File" (page 62)
- "Starting and Stopping xntpd" (page 69)
- "Using ntpq to Query Systems Running xntpd" (page 70)
- "Troubleshooting NTP" (page 72)

Stratum Levels and Time Server Hierarchy

An NTP **synchronization subnet** is a network of time-keeping systems, called **time servers**. These time servers are a subset of the systems on a network or an internetwork. Each time server synchronizes its time with the the Universal Coordinated Time (UTC), and measures the time difference between its local system clock and the neighboring system clocks. These servers are automatically assigned stratum values, which indicate how close the time server is to the time source.

Stratum-1 Time Servers

Time servers are organized into levels, or **strata**. Stratum-1 servers are directly connected to an external time source. The stratum-1 server relies on the external source of time to provide the correct time, and synchronizes its system clock to that external time source. The external time source can be a device such as a radio receiver. Figure 4-2 shows the relationship between the GPS receiver time source and the stratum-1 server associated with it.

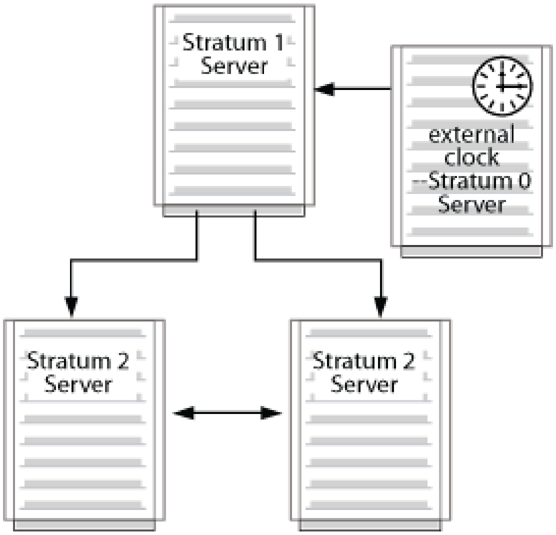

Figure 4-2 Stratum-1 Time Servers

Stratum-2 and -3 Time Servers

Stratum-2 time servers use stratum-1 servers as their time source. Likewise, stratum-3 servers use stratum-2 servers as their time sources. The maximum stratum level a server can have is 15.

Time Server Roles

An NTP time server can take different roles in its relationships with other time servers in the synchronization subnet. A time server can take one or more of the following roles:

- **Server**— Provides time to clients when requested. This role can be assumed by time servers at various strata.
- **Peer**— Obtains time from a specified server and provides time to that server, if requested. This role is most appropriate for stratum-1 and stratum-2 servers.
- **Client**— Obtains time from a specified server, but does not provide time to that server. This role is appropriate for time servers that obtain time from a server of a lower-numbered stratum (for example, a stratum-1 server). The local host can in turn provide synchronization to its clients or peers.
- **Broadcaster**— Provides time to the specified remote host, or more typically, the broadcast address on a LAN. This role is most appropriate for an NTP time server that provides time to workstation clients on a LAN.
- **Broadcast Client**— Listens for and synchronizes with the broadcast time. This role is most appropriate for time server clients on a LAN.

NOTE: Broadcasting is not recommended especially when used with local clock impersonators. Broadcasting is an old concept that is no longer used.

Figure 4-3 illustrates the relationship between time servers in a synchronized subnet.

Figure 4-3 Example of Relationships Between Time Servers

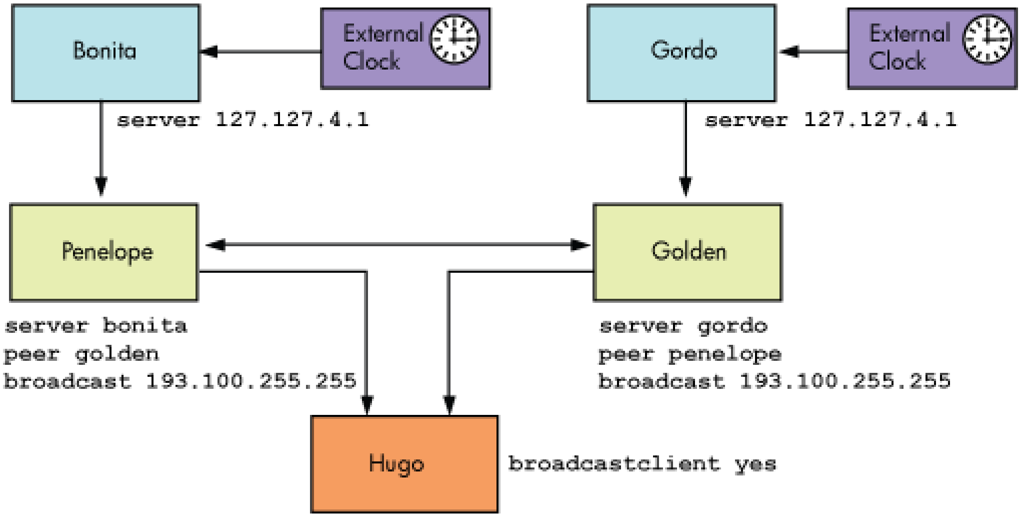

Planning a Multiple-Server NTP Configuration

You must consider the following guidelines when planning your configuration:

- Every NTP hierarchy must have atleast one stratum-1 server. You can configure the administrative domain to contain outside sources of synchronization, which ultimately link to stratum-1 server, or you can implement your own hierarchy of NTP time servers with one or more stratum-1 servers.
- Configure atleast three time servers in the administrative domain because it is important to provide multiple, redundant sources of time synchronization. NTP is designed to select an optimal source of synchronization from several sources. Each time server must be a peer with each of the other time servers.
- For each time server, select one or three *outside* sources of synchronization. This assures a relative degree of reliability in obtaining time, when you select sources that do not share common paths. The sources must operate at a stratum level that is one less than the local time servers.
- Each outside source of synchronization must be in different administrative domains, and you must access the sources from different gateways and access paths. You must avoid loops and common points of failure. Do not synchronize multiple time servers in an administrative domain to the same outside source.

- For enterprise networks that contain many file servers and workstations, the local time servers must obtain service from stratum-1 servers.
- While defining a relationship between a higher-numbered stratum server and lower-numbered stratum server, configure the relationship in the higher-numbered stratum server. For example, if a stratum-3 server is a client of a stratum-2 server, configure the relationship in the stratum-3 server. This simplifies configuration maintenance, because configuration of the higher-numbered stratum server changes often.

Configuring NTP using the Configuration File

This section describes the statements that you can define in the `/etc/ntp.conf` configuration file. It discusses the following topics:

- "Configuring Relationships with Other Time Servers" (page 62)
- "Configuring an External Clock" (page 64)
- "Configuring a Driftfile" (page 65)
- "Configuring Authentication" (page 66)

Configuring Relationships with Other Time Servers

The role of a time server depends on its relationship with other servers in the synchronization subnet. In the configuration file, you can define a role using one of the following statements for `peer`, `server`, `broadcast` and `broadcastclient`:

- `peer` host | IP_address specifies that host must provide time to the local host with which the local host can synchronize its time, and the local host must also provide time to the host.
- `server` *`host`* `|` *`IP_address`* specifies that host must provide time that the local host can synchronize to, and the local host does not provide time to which the host can synchronize to. (The local host is a client for the host.) Additionally, you can define server statements to configure external clocks (radio clocks or local system clocks) for stratum-1 servers. See "Configuring an External Clock" (page 64) for more information.
- `broadcast` *`host`* `|` *`broadcast_address`* specifies that the `xntpd` daemon in the local host transmits broadcast NTP messages to the broadcast_address, usually the broadcast address on the local network (The local host is a broadcaster.)

You can specify one or more of the following options using the `peer`, `server`, or `broadcast` statements:

- `key`*`number`*

  This option specifies that the NTP packets sent to the host are encrypted using the key that is associated with *`number`*. You must enable the authentication feature of `xntpd` for this option. See "Configuring Authentication" (page 66) for more information.

- `version 1`

  You must specify this option if `xntpd` requests time from a host running `ntpd`, a daemon that is based on version 1 of the NTP protocol. You must specify version 2 if `xntpd` requests time from a host running an `xntpd` implementation that is based on version 2 of the NTP protocol. If you do not specify either of these options, `xntpd` daemon sends version 3 NTP packets when polling the host. If the host contains a version 1 or version 2 NTP protocol, the packets are discarded.

- `prefer`

  This option specifies that the host must be the primary source for synchronization when it is one of several valid sources. This option is useful for a time server on a high-speed LAN that is equipped with an external time source, such as a radio clock. You can use external sources for time synchronization. However, the local time server must be the preferred synchronization source.

Apart from these options, you can define the broadcast client in the configuration file using the `broadcastclient` yes | no statement. The statement `broadcastclient` yes indicates that the local host must listen for and attempt to synchronize with the broadcasted NTP packets. The optional statement `broadcastdelay` *`seconds`* specifies the default round trip delay for a broadcaster.

NOTE: Every node in an NTP hierarchy must have either a `server` statement or a `broadcastclient yes` statement in its configuration file. Every node must have an upper-level server. A stratum-1 server must also have a `server` statement in its configuration file, which specifies a radio clock or internal system clock as a time source.

If the local host assumes the role of a server in providing time to the clients, you need not configure the local host as a time server on the local system. Instead, you must specify the local host name or IP address in the `server` statement in the client system's configuration file.

If authentication is enabled on the local host, the roles you configure are subject to the authentication process. For example, if you configure the local host as a peer or a client of a stratum-1 server, and the remote server does not satisfy the standards for an authenticated synchronization source, the local host does not use the remote server as a time source. See "Configuring Authentication" (page 66) for more information.

NOTE: `xntpd` is an HP implementation of version 3.2 of a publicly-available NTP daemon. HP does not guarantee that `xntpd` is fully compatible with version 1 or version 2 implementations of the daemon.

Configuring an External Clock

Clocks are normally configured with `server` statements in the configuration file. You can configure `xntpd` to support an external clock. You can insert the clock address anywhere in the configuration file.

Clocks are referenced by an address of the format `127.127.`*`t`*`.`*`u`*, where *`t`* specifies the clock type and *`u`* specifies the unit number, which depends on the clock type for interpretation (this allows multiple instances of the same clock type on a single host).

`xntpd` supports the following clocks:

- Netclock/2 WWVB Synchronized Clock.

  A stratum-1 server is configured with this type of clock. The address used to configure the clock is `127.127.4.`*`u`*, where *`u`* is a value between 1 and 4. You must also create a device file `/dev/wwvb%`*`u`*.

- Local Synchronization Clock (pseudo clock)

  A system with this type of clock uses the local system clock as a time source. The address used to configure this clock is `127.127.1.`*`u`*, where *`u`* is a value between 0 and 15 and specifies the stratum level at which the clock runs. When you synchronize the local host to this clock, the local host operates at one stratum level higher than the local clock. You can use this type of clock in an isolated synchronization subnet, which does not have access to a stratum-1 time server.

For more information on configuring external clocks, type `man 1M xntpd` at the HP-UX prompt.

Figure 4-4 shows the peer, server, and broadcast statements that are configured for all the servers.

Figure 4-4 Example Configurations

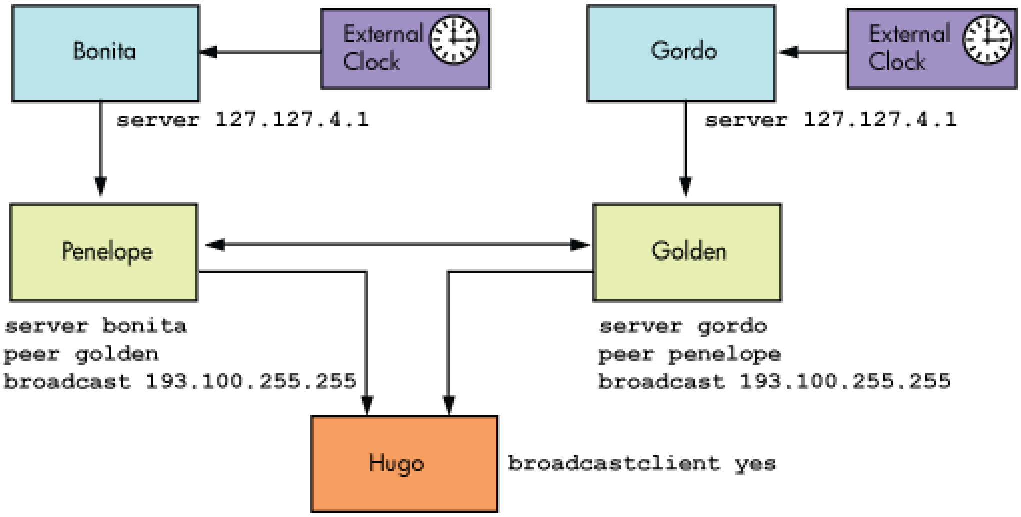

You must configure the time server in the client system. For example, if `Penelope` is a client for `Bonita`, you must configure the name or IP address of `Bonita` on `Penelope`. You need not configure `Penelope` as a client on `Bonita`.

Configuring a Driftfile

The `xntpd` daemon computes the clock frequency error for a local host, and stores the frequency error in a driftfile. `xntpd` computes an accurate estimate of the frequency error after running for one or more days. When `xntpd` restarts, the driftfile allows `xntpd` to reinitiate itself to the estimate stored in the driftfile; thus, helping `xntpd` to save time in recomputing the frequency estimate. You can specify a name and path for the driftfile.

NOTE: `xntpd` must be running continuously; if you wish to stop `xntpd`, it must be for a short duration.

You can use the following option to specify the driftfile:

```
driftfile driftfile
```

where, driftfile specifies the file name used to record the frequency offset of the local clock oscillator. HP recommends the location `/etc/ntp.drift` for storing the driftfile. The following is an example of a `driftfile` statement:

```
driftfile /etc/ntp.drift
```

Authentication is a mechanism used to prevent unauthorized access to time servers. Authentication is enabled on a system-by-system basis. Once enabled on a system, authentication applies to *all* NTP relationships configured on the system. If you enable authentication on a host, the host synchronizes time only with those time servers that send messages encrypted with a configured key.

In an authenticated mode, each NTP packet transmitted by a host is appended by a **key number** and an **encrypted checksum** of the packet contents. The key number is specified in the `peer`, `server`, or `broadcast` statement for the remote host. You specify use the Data Encryption Standard (DES) or the Message Digest (MD5) algorithm to encrypt the NTP packets.

Upon receipt of an encrypted NTP packet, the receiving host recomputes the checksum and compares it with the checksum included in the packet. Both, the sending and receiving systems must use the same encryption key defined by the key number.

When authentication is enabled on a host, the host does not consider the following time servers for synchronization:

- Time servers that send unauthenticated NTP packets.
- Time servers that send authenticated packets that the host is unable to decrypt.
- Time servers that send authenticated packets encrypted using a non-trusted key.

An **authentication key file** is specified on the host and contains a list of keys and their corresponding key numbers. Each key-key number pair is further defined by a key format, which determines the encryption method. For more information about the authentication key file, type `man 1M xntpd` at the HP-UX prompt. A sample key file is provided in `/usr/newconfig/etc/ntp.keys`. HP recommends the location `/etc/ntp.keys` for storing the key file. You must secure the key file by giving the permission 600.

While the key file can contain many keys, you can declare a subset of these keys as **trusted keys**. Trusted keys are used to determine if a time server is trusted as a potential synchronization candidate. Only time servers that use a specified trusted key for encryption, and whose authenticity is verified by successful decryption, are considered synchronization candidates.

Figure 4-5 illustrates how authentication works.

Figure 4-5 Authentication Example

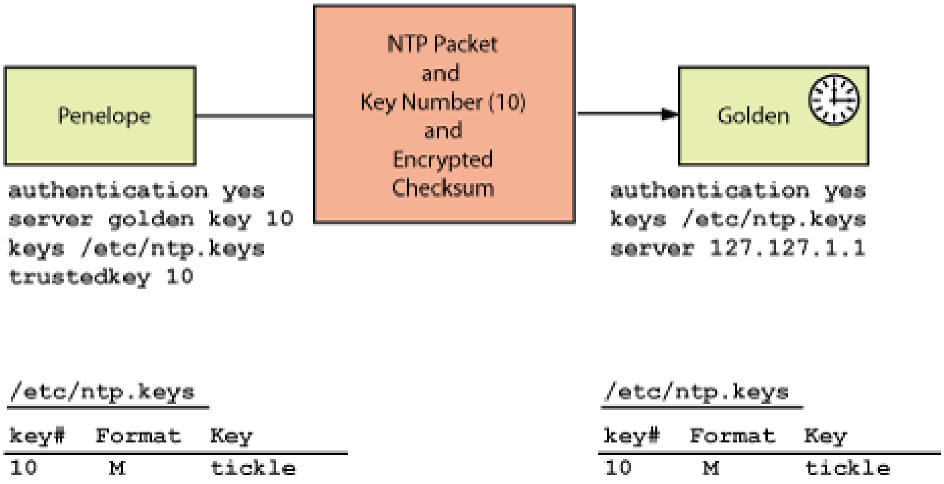

```
authentication yes
server golden key 10
keys /etc/ntp.keys
trustedkey 10
```

```
authentication yes
keys /etc/ntp.keys
server 127.127.1.1
```

/etc/ntp.keys

key#	Format	Key
10	M	tickle

/etc/ntp.keys

key#	Format	Key
10	M	tickle

In Figure 4-5, authentication is enabled for both `Penelope` and `Golden`. An NTP time request from `Penelope` to `Golden` includes the authentication fields – key ID (10), and a checksum, tickle, encrypted with the key corresponding to the key ID 10. When `Golden` receives this request, it recomputes the checksum using the packet's key ID field (10) to look up for the key ID 10 in its key file (tickle) and compares the checksum with the authentication field in the request.

`Golden` sends back time information with the key ID 10 and a checksum encrypted using the encryption key tickle.

Additionally, `Penelope` accepts time synchronizations from hosts that have used the key ID 10 and the corresponding encryption key tickle.

To enable authentication on the local host, include the following statement in the `/etc/ntp.conf` configuration file:

```
authenticate yes
```

If you do not specify this statement, authentication is not enabled. When you enable authentication, you can specify the following options:

- `-e authdelay`

  This option indicates the amount of time (in seconds) required to encrypt an NTP authentication field on the local host.

IMPORTANT: The startup script automatically calculates the proper value for `authdelay` for the local system and writes it into the configuration file `/etc/ntp.conf`. Do *not* modify this value.

- `-k keyfile`

  This option specifies the file that contains the encryption keys used by `xntpd`.

- `-t key`

  This option specifies the encryption key IDs that are trusted as synchronization sources.

Restricting Incoming NTP Packets

`xntpd` provides a mechanism for restricting access to the local daemon from certain sources. In the `/etc/ntp.conf` file, you can define a **restriction list** that contains the addresses or addresses and masks of sources that may send NTP packets to the local host. For each address or address-mask specified in the restriction list, you can define flags to restrict time service or queries to the local host.

The source address of each incoming NTP packet is then compared to the restriction list. If a source address matches an entry in the restriction list, the restriction defined by the corresponding flag is applied to the incoming packet. If an address-mask is specified in the restriction list, the source address of each incoming NTP packet is ANDed with the mask, and then compared with the associated address for a match.

The restriction list should not be considered an alternative to authentication. It is most useful for keeping unwanted or broken remote time servers from affecting your local host. An entry in the restriction list has the following format:

```
restrict address [mask mask] [ntpport] [flag] [flag2]...
```

The keyword `ntpport` causes the restriction list entry to be matched only if the source port in the packet is the NTP UDP port 123.

Table 4-6 shows the flags that can be specified for `xntpd`:

Table 4-6 Restrict Option Flags

Flag	Effect
`ignore`	Ignore all packets.
`noquery`	Ignore `ntpq` queries.
`nomodify`	Ignore `ntpq` packets that attempt to modify the state of the server.
`noserve`	Ignore requests for time, but permit `ntpq` queries.

Table 4-6 Restrict Option Flags *(continued)*

Flag	Effect
`nopeer`	Provide time service, but do not form peer association.
`notrust`	Do not use the host as a synchronization source.

A restriction list entry with no flags set leaves matching hosts unrestricted. A source address of an incoming packet may match several entries in the restriction list. The entry that matches the source address most specifically is the entry that is applied. For example, consider the following restriction list entries:

```
restrict 193.100.0.0 mask 255.255.0.0 ignore restrict 193.100.10.8
```

The first entry causes packets from source addresses on net 193.100 to be ignored. However, packets from host 193.100.10.8 are unrestricted, as specified by the second entry. The two restriction list entries effectively cause all packets from net 193.100 to be ignored, with the exception of packets from host 193.100.10.8.

The following are examples of restriction list entries for a local host with the address 193.100.100.7. These entries assume that `ntpq` requests to the local host can be made only from the local host or the host with address 193.8.10.1, while the local host only synchronizes to a time source on net 193.100.

```
#default entry - matches *all* source addresses restrict default notrust nomodify
#trust for time, but do not allow ntpq requests restrict 193.100.0.0 mask 255.255.0.0 nomodify noquery
#ignore time requests, but allow ntpq requests restrict 193.8.10.1 noserve  #local host address is unrestricted
restrict 193.100.100.7
```

Starting and Stopping xntpd

To start `xntpd`, do one of the following:

- Set the environment variable `XNTPD` to `1` in the file `/etc/rc.config.d/netdaemons`. This causes `xntpd` to start automatically when you boot the system.
- Issue the following command to run the `xntpd` startup script:

  ```
  /sbin/init.d/xntpd start
  ```

You can specify command-line arguments for starting `xntpd` with the `XNTPD_ARGS` environment variable in the file `/etc/rc.config.d/netdaemons`. For more information on command-line arguments, type `man 1M xntpd` manpage at the HP-UX prompt.

NOTE: xntpd must be running continuously; if you wish to stop xntpd, it must be for a short duration.

If you modify the configuration file or the XNTPD_ARGS environment variable in the file /etc/rc.config.d/netdaemons while xntpd is running, you have to stop and restart the daemon for the configuration changes to be effective.

To stop xntpd, issue the following command:

```
/sbin/init.d/xntpd stop
```

Using ntpq to Query Systems Running xntpd

The standard network time protocol query program, ntpq, is used to query systems that implement the NTP mode 6 control message, about the current state of the server. It can also be used to obtain a list of a server's peers. ntpq sends requests to and receives responses from NTP time servers using a special form of NTP messages called mode-6 control messages. You can run ntpq either in the interactive mode or from the command line.

ntpq is useful for querying remote NTP implementations to assess their timekeeping accuracy and to expose problems in configuration or operation. For more information, type man 1M ntpq at the HP-UX prompt.

NOTE: When you specify time-related configuration options in the /etc/ntp.conf file, you specify the values in seconds. ntpq, however, displays time values in milliseconds, as specified by RFC 1305 (Network Time Protocol (Version 3) – Specification and Implementation) NTP standard.

Verifying ntpq

Use ntpq to verify whether:

- xntpd can form associations with other NTP hosts.
- Synchronization is happening correctly.

After xntpd starts, run the ntpq program with the -p option:

```
/usr/sbin/ntpq -p
```

The ntpq Program Output

The -p option prints a list of peers known to the server, along with a summary of their states as shown in Table 4-7.

Table 4-7 An ntpq Output Indicating Known NTP Hosts

```
remote      refid     st t  when   poll reach  delay  offset  disp
==================================================================
*GPS_HP(1)  GPS       0  l    48    64  377     0.00    0.516   4.19
hpps.cup.hp cupertino 3  u   467  1024  377     7.20  -12.430  15.67
+server2     WWVB      1  u   173    256  377   279.95   20.56   16.40
+node1      node3      2  u   131    256  373     9.89   16.28   23.25
```

Each column in Table 4-7 is described as follows:

- The `remote` (server name) column depicts the hosts specified in the local host's configuration file plus other hosts that are configured as peers with the local host. The host address can be preceded by the following special characters:
  - □ `*` indicates the current synchronization source.
  - □ `#` indicates that the host is selected for synchronization, but distance from the host to the server exceeds the maximum value.
  - □ `o` indicates that the host is selected for synchronization, and the PPS signal is in use.
  - □ `+` indicates the host included in the final synchronization selection set.
  - □ `x` indicates that the host is the designated false ticker by the intersection algorithm.
  - □ `.` indicates that the host is selected from the end of the candidate list.
  - □ `-` indicates a host discarded by the clustering algorithm.
  - □ `blank` indicates a host is discarded due to high stratum and/or failed sanity checks.
- The `refid` (reference identification) column indicates the current source of synchronization for the remote host. `.WWVB.` indicates that the host uses a radio clock that receives time signals from the U.S. government radio station `WWVB`.
- The `st` (stratum) column indicates the stratum level of the remote host.
- The `t`(types) column denotes the available types, which include
  - — `l`=local (such as a GPS clock)
  - — `u`=unicast (this is the common type)
  - — `m` = multicast
  - — `b`= broadcast
  - — `-` = netaddr (usually 0)
- The `when` column indicates the number of seconds since the remote host response was received.
- The `poll` (poll period) column indicates the polling interval to the remote host, as determined by `xntpd`. You can define the minimum polling interval with the `minpoll` option in the `peer`, `server`, or `broadcast` definitions in the `/etc/ntp.conf` file. Some popular values for network connections include 512

and 1024 seconds (approximately 8 mins. and 17 mins.). Systems with external clocks, like GPS, must poll every 64 seconds or less.

- The `reach` (reachability) column indicates how successful attempts to reach the server are. This is an 8-bit shift register with the most recent probe in the 2^0 position. The value 001 indicates the most recent probe was answered, while 357 indicates one probe was not answered. The value 377 indicates that all the recent probes have been answered.
- The `delay` (round trip time) column indicates the time (in milliseconds) taken by the reply packet to return in response, to a query sent by the server.
- The `offset` (time difference) column indicates the time difference (in milliseconds) between the server's clock and the client's clock. When this number exceeds 128, and the message `synchronization lost` appears in the log file.
- The `disp` (dispersion) column indicates the difference in the `offset` measurement between two samples. This is an error-bound estimate. The dispersion is a primary measure of the network service quality.

Troubleshooting NTP

This section outlines techniques that can help you diagnose and correct common problems with the NTP.

Verifying That xntpd is Running

Issue the following command to determine out if `xntpd` is running:

```
/usr/bin/ps -ef | /usr/bin/grep xntpd
```

This command reports the process identification (PID), current time, and the command invoked (`xntpd`). Following is an example output:

```
daemon  4484  1  0  Feb 18   ?    0:00    xntpd
```

Ensure that `syslogd` is configured to log daemon information messages to the file `/var/adm/syslog/syslog.log`. To check this configuration, ensure that the file `/etc/syslog.conf` includes one of the following entries:

```
*.info    /var/adm/syslog/syslog.log
```

or

```
daemon.info   /var/adm/syslog/syslog.log
```

If `xntpd` is not running, check the `syslog` file for related messages.

NTP Associations

Each NTP daemon must form an association with a time source: a higher-level (lower stratum) server for stratum-1 servers, or an external clock. NTP daemons can form additional associations with peer servers. Use the following command to list the NTP associations established by the local NTP daemon:

```
/usr/sbin/ntpq -p
```

In the output, an asterisk (*) must appear next to the node name to indicate that an association has been formed.

Table 4-8 indicates that the local NTP daemon has established an association with the NTP daemon on node `good.cup.hp`.

Table 4-8 An ntpq Output Indicating NTP Associations

```
remote        refid      st    when  poll   reach   delay    offset   disp
======================================================================
*good.cup.hp  LOCAL(1)   2     29    64     377     5.43     -0.16    16.40
bad           0.0.0.0    -     31    64     0
```

If the local node is unable to form an association with its higher-level server or its peer, you must login to the higher-level server or peer and issue the following command:

```
/usr/sbin/ntpq -p
```

Verify whether the higher-level server or peer has established an association with a time source.

Query with Debug Option

If you cannot form an association with a server or peer, stop the local `xntpd` and send a time request to the server or peer using the `ntpdate` command and the debug (`-d`) option, as shown in the following example:

```
#/sbin/init.d/xntpd stop

#/usr/sbin/ntpdate -d server
```

The debug (`-d`) option prints information about the requests sent to the remote `xntpd` daemon, and the information returned by the remote `xntpd`. The `ntpdate` command fails if `xntpd` is already running on the local system. Also, the `ntpdate` command does not use authentication; therefore, it must be executable only by the `root`.

You can also use `ntpdate` on systems where exact time synchronization is not necessary. You can run `ntpdate` periodically from `cron` to synchronize the local clock with the other system's clock. For more information, type `man 1M ntpdate` at the HP-UX prompt.

Error Messages

This section describes the error messages that you may encounter while working with NTP.

No server suitable for synchronization found.

This message indicates that the NTP server is not responding. Packets were sent out, but a reply was not returned. The reason may be that the server is down, or the network link is broken or extremely congested. Or, perhaps the NTP daemon died on the server and has not locked on to its time sources. NTP version 3.5 and above take between 5 and 15 minutes after starting up the daemon to synchronize, and it does not respond to client requests during this time.

Last adjustment did not complete.

This message indicates that NTP is attempting to make adjustments, larger than the system's maximum slew rate allows, in one clock tick. Therefore, the remaining adjustments are pushed to the next clock tick. This is handled automatically. You can notice this message during the first hour after the NTP daemon is started. If this message continuous to appear after a few days of steady operation, this indicates that your system clock is drifting. This may result in loss of contact with the network time server.

Synchronization lost.

This message indicates that NTP has cleared the statistics registers, and has started evaluating the available time servers to choose the best time server. This message appears when a step adjustment (greater than 128 milliseconds) is done because the step leaves the system unsynchronized by definition. If the system does many step adjustments, it indicates a network congestion problem. To review this problem, do the following steps:

1. Run **`ntpq -p`**
2. Examine the dispersion statistics.

Common Problems

This section covers typical problems with `ntp` operation.

Problem 1: No suitable server for synchronization found.

Every NTP time hierarchy must have at least one stratum-1 server configured with an external time source, such as, an attached radio clock (Netclock/2 WWVB Synchronized Clock) or the local system clock. If a stratum-1 server in the hierarchy does not exist, association is not formed. To verify whether the local `xntpd` is able to form an association, issue the following command:

```
/usr/sbin/ntpdate server
```

The *server* is the name of a trusted server, such as a peer or high-level (lower stratum) server. If the local xntpd is unable to form any association, this command returns the message No suitable server for synchronization found. The possible causes for this error message is discussed in the following sections.

Time Difference Greater than 1000 seconds

While evaluating incoming time updates, clients and peers reject time from servers or peers if the time difference is greater than 1000 seconds. If xntpd is configured with the broadcastclient entry in the /etc/ntp.conf file, xntpd dies if it cannot find a suitable server after six consecutive polls, or five polling cycles (approximately 320 seconds if using the default polling interval). If xntpd is configured with server entries in the /etc/ntp.conf file, xntpd dies if it cannot find a suitable server after five consecutive polls, or four polling cycles. The number of polls for the broadcastclient entry configuration is one count more than the server entry configuration because the first successful reception of an NTP message for the broadcastclient entry adds the timeserver to the list of timeservers.

Because of this behavior, you may have to issue the following command to synchronize the local system time with another NTP server before starting xntpd:

```
/usr/sbin/ntpdate server
```

When a server entry is configured in the /etc/ntp.conf file, xntpd knows the name of the timeserve and need to wait for only five polls to synchronize with the timeserver.

The disp value in the ntpq -p command starts at approximately 16000 seconds and is divided by two when a successful NTP message is received from the timeserver. If a value lower than 1000 seconds is received, the client or peer attempts to synchronize with the timeserver and an error message is received about the time difference. The first reception writes a value of 16000 seconds to the disp value for the broadcastclient entry. Because the client or peer polls the timeserver in case of a server entry, the original disp value is set to disp.

For HP-UX NFS Diskless Clusters, the /sbin/init.d/xntpd script on the diskless clients executes xntpdate to synchronize time with the diskless cluster server before starting xntpd.

You can also specify a trusted time server explicitly in the file /etc/rc.config.d/netdaemons, and /sbin/init.d/xntpd will execute xntpdate, querying the specified time server.

Startup Delay

When xntpd is started, it takes five poll cycles (320 seconds using the default polling interval) to form an association with a higher-level server or peer. During this time window, xntpd does not respond to time requests from other NTP systems, because it does not have a suitable time source. This window exists even though xntpd is using an external clock, which can be either an attached radio clock (Netclock/2 WWVB Synchronized Clock) or the local system clock (server 127.127.*n*.*n*).

For external clocks, xntpd does not form a complete association until it has sent five successful polls to itself using the local loopback address.

Problem 2: Version 1 and 2 NTP Servers Do Not Respond

NTP version 3 packets are ignored by NTP version 1 and version 2 systems. The solution is to indicate the version 1 and 2 system in the configuration entry on the version 3 system. This informs the version 3 system to use the older message formats when communicating with these systems.

The following configuration file entries inform xntpd to use NTP version 2 message formats when communicating with some_ver2.sys and NTP version 1 when communicating with some_ver1.sys.

```
server some_ver2.sys version 2

server some_ver1.sys version 1
```

Reporting Problems

Provide the following information while reporting NTP problems:

- The configuration file /etc/ntp.conf (or an alternate configuration file)
- The /etc/rc.config.d file
- NTP driftfile (if configured)
- NTP statistics file (if configured)
- The /var/adm/syslog/syslog.log file (xntpd/NTP entries)
- The /usr/sbin/ntpq -p output
- The ntpdate -d server output (stop the local xntpd first).

5 Troubleshooting Internet Services

This chapter describes how to troubleshoot the Internet Services software.

It discusses the following topics:

- “Troubleshooting Overview” (page 77)
- “Troubleshooting Tips” (page 80)
- “Reporting Problems to Your Hewlett-Packard Support Contact” (page 91)

Troubleshooting Overview

Troubleshooting data communications problems may require you to investigate many hardware and software components. Some problems can be quickly identified and resolved. These include invalid software installation, version incompatibilities, insufficient HP-UX resources, corrupt configuration shell scripts, and programming or command errors. Other problems require more investigation.

Once identified, most problems can be resolved by the programmer, user, or node manager, using the suggestions in this chapter or the error messages documented in the link installation manuals. However, there may be problems that you should report to your Hewlett-Packard support contact. This chapter includes guidelines for submitting an HP Service Request (SR).

This section discusses the following topics:

- “Characterizing a Problem” (page 77)
- “Diagnostic Tools Summary” (page 78)
- “Diagnosing Repeater and Gateway Problems” (page 79)

Characterizing a Problem

It is important that you ask questions when trying to characterize a problem. Start with global questions and gradually get more specific. Depending on the response, ask another series of questions until you have enough information to understand exactly what happened. Key questions to ask are as follows:

- Does the problem seem isolated to one user or program? Can the problem be reproduced? Did the problem occur under any of the following circumstances:
  — When running a program?
  — When issuing a command?

- — When using a nodal management utility?
  - — When transmitting data?
- Does the problem affect all users? The entire node? Has anything changed recently? The possibilities are as follows:
  - — New software and hardware installation.
  - — Same hardware but changes to the software. Has the configuration file been modified? Has the HP-UX configuration been changed?
  - — Same software but changes to the hardware. Do you suspect hardware or software?

It is often difficult to determine whether the problem is hardware-related or software-related. The symptoms of the problem that indicate you should suspect the hardware are as follows:

- Intermittent errors.
- Network-wide problems after no change in software.
- Link-level errors, from logging subsystem, logged to the console.
- Data corruption—link-level trace that shows that data is sent without error but is corrupt or lost at the receiver.
- Red light on the LAN card is lit, or yellow light on the X.25/800 card is lit.

Following are symptoms that would lead you to suspect the software:

- Network services errors returned to users or programs.
- Data corruption.
- Logging messages at the console.

Knowing what has recently changed on your network may also indicate whether the problem is software-related or hardware-related.

Diagnostic Tools Summary

Table 5-1 describes the most frequently used diagnostic tools.

Table 5-1 Diagnostic Tools

Tool	Description
`netstat`	A nodal management command that returns statistical information regarding your network.
`landiag`	A diagnostic program that tests LAN connections between HP computers.

Table 5-1 Diagnostic Tools *(continued)*

Tool	Description
`nwmgr`	A diagnostic program that runs link-level loopback tests between HP systems. `nwmgr` uses IEEE 802.3 link-level test frames to check physical connectivity with the LAN. This diagnostic tool is different from the loopback capability of `landiag` because it tests only the link-level connectivity and not the transport-level connectivity.
`ping`	A diagnostic program that verifies the physical connection to a remote host and reports the round-trip communication time between the local and remote hosts. (Type `man 1M ping` for more information.)
`psidad`	A utility under DUI that can help to identify problems on the PSI/800 board/card.
`rlb`	A diagnostic program that tests LAN connections to other HP computers. `rlb` does not test a connection to an HP 1000 computer.
`x25check x25server`	These two work in tandem. `x25server` runs on the logically remote host (could be same physical host) and echoes packets sent to it over the X.25 network by `x25check`.
`x25stat`	A nodal management command that returns status and information from the X.25 device and card. It provides interface status configuration information and virtual circuit statistics.
`x25upload`	A command that uploads the firmware in case of problems with the firmware on the board.
Event Logging	A utility that sends informational messages regarding network activity to the system console or to a file.
Network Tracing	A utility that traces link-level traffic to and from a node. HP recommends that you enable tracing only when troubleshooting a problem unsolved by other means.

Diagnosing Repeater and Gateway Problems

If you are using a repeater, and hosts on either side of the repeater are having difficulty communicating with each other, a repeater subsystem failure may have occurred. In Figure 5-1, all of the systems on side A are able to communicate with each another. All the systems on side B are able to communicate with each other. If communication is cut from side A to side B, the repeater subsystem is suspect for causing the fault, because it is the medium by which side A and side B communicate.

Figure 5-1 Troubleshooting Networks that Use Repeaters

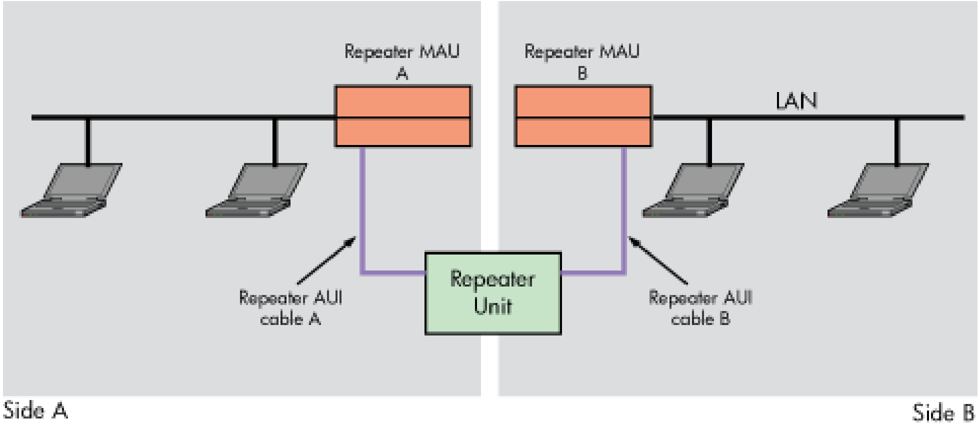

The same concept holds for communication through a gateway. If you suspect a gateway problem, try the following procedures:

- To determine if you are set up to communicate with the desired node, execute the following:

  ```
  netstat -r
  ```

- To obtain routing statistics, execute the following:

  ```
  netstat -rs
  ```

The statistics could indicate a bad route, suggesting a problem with a gateway node. To identify such errors, do the following:

- Check with the node manager of the gateway node to ascertain proper operation of the gateway.
- You can detect problems with the X.25 line by the number of errors shown when you execute the following:

  ```
  x25stat -f -d /devicefile
  ```

For more information on troubleshooting gateways, see the appropriate link manual. For information on repeaters, see the *HP-PB LAN Interface Controller (LANIC) Installation Manual*.

Troubleshooting Tips

This section provides useful tips for troubleshooting the Internet Services software.

When troubleshooting problems with the Internet Services, you need a reference point to work from. For example, does the problem exist on the remote system or on the local system? However, the terms local and remote are limited in their description of complex communications, such as when a local system logs on to a remote system and then the remote system logs back on to the local system. At that point, which is the local system and which is the remote system?

A better solution is to use the terms client and server. The term client refers to a process that is requesting a service from another process. The term server refers to a process or host that performs operations requested by local or remote hosts that are running client processes.

HP has implemented a super-server known as the Internet daemon, `inetd`. This program acts like a switchboard; that is, it listens for any request and activates the appropriate server based on the request.

A typical network service consists of two co-operating programs. The client program runs on the requesting system. The server program runs on the system with which you want your system to communicate. The client program initiates requests to communicate. The server program accepts requests for communication. For example, the network service `rlogin` is the client program that requests a login to a remote HP-UX or other UNIX system. When the request to log in is received on the remote host by `inetd`, `inetd` invokes the server program for `rlogin` (called `rlogind`) to handle the service request.

Flowchart Format

The flowcharts in this section each have a corresponding set of labeled explanations. You can follow the flowcharts alone or follow the flowcharts and read the explanations for more detail. The explanations are on the pages that follow each flowchart.

Figure 5-2 Flowchart Symbols

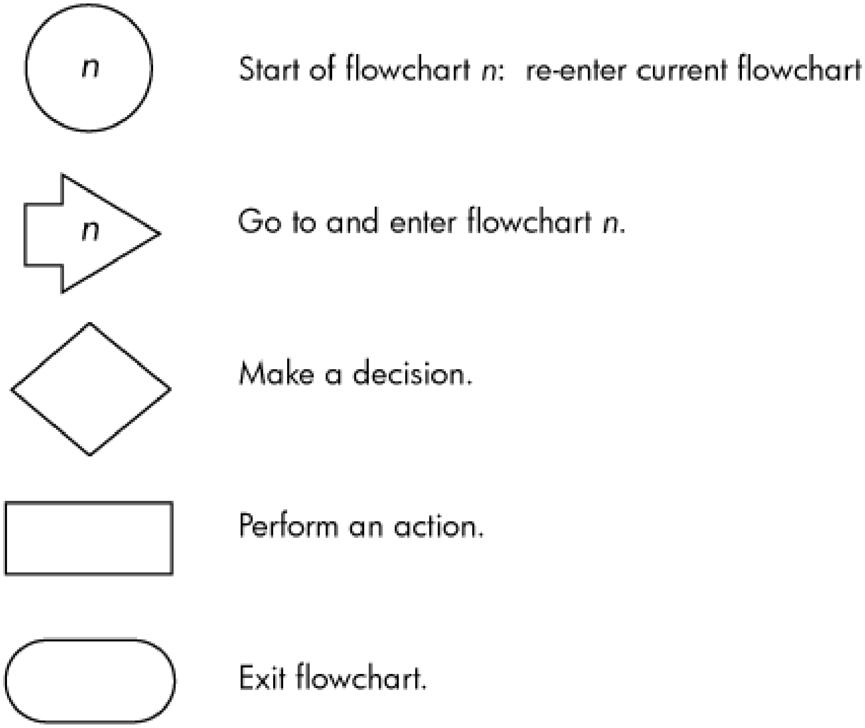
n
Start of flowchart *n*: re-enter current flowchart
n
Go to and enter flowchart *n*.
Make a decision.
Perform an action.
Exit flowchart.

Error Messages

The error messages generated by a service as seen on the client can be generated by the client or by the server. Error messages from the client occur before a connection is completely established. Error messages from the server occur after a connection is completely established.

Whenever you receive an error message, follow the corrective action supplied in the manpage for that service. The error message is preceded by the name of the service. Table 5-2 shows the appropriate manpage to consult for a description of the error messages.

Table 5-2 Manpages for Error Messages

Service	Client	Server
`telnet`	`telnet`(1)	`telnetd`(1M)
`ftp`	`ftp`(1)	`ftpd`(1M)
`rlogin`	`rlogin`(1)	`rlogind`(1M)
`remsh`	`remsh`(1)	`remshd`(1M)
`rcp`	`rcp`(1)	`remshd`(1M)

Table 5-2 Manpages for Error Messages *(continued)*

`ruptime`	`ruptime`(1)	`rwhod`(1M)
`rwho`	`rwho`(1)	`rwhod`(1M)
`ddfa`	user application	`ocd`(1M)

If the server or the client is not an HP system, see the appropriate user's manual or system administration manual for that system. There is not a standard naming convention for servers or processes that activate the servers; however, you should be able to find the information in the system's documentation.

Services Checklist

Following is a services checklist to help you diagnose the problem:

- Answer the questions in the section "Characterizing a Problem" (page 77) at the beginning of this chapter.
- Run the service to your own node. To do this, your node name and Internet address must be in the `/etc/hosts` file. If the server is successful, then the client and the server halves of the service operate correctly. This provides a starting point to determine where problems are occurring.
- Use the flowcharts that follow to help you identify what is causing the problem.

Flowchart 1. Checking for a Server

Follow Flowchart 1 for all services and servers, and replace the words service and server with the appropriate service name or server name.

Figure 5-3 Flowchart 1. Checking for a Server

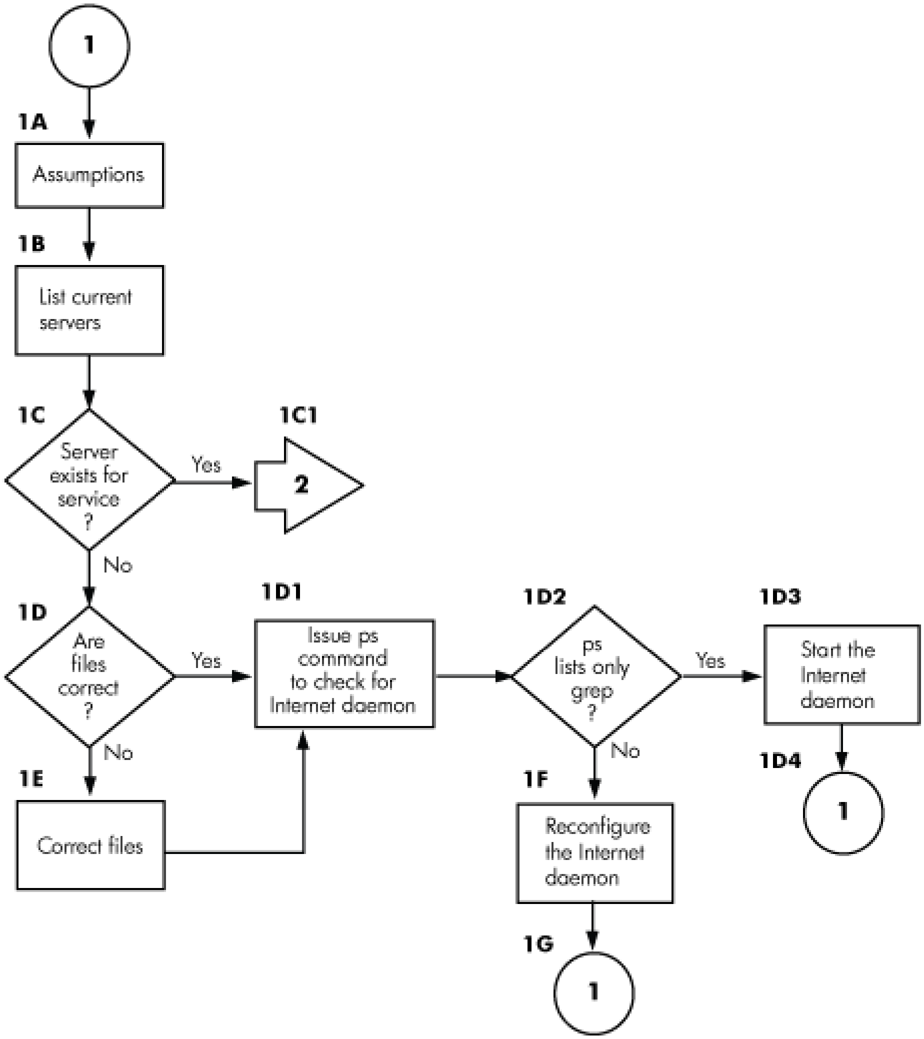

1A. Assumptions. Before you begin Flowchart 1, you should have verified local node operations and verified connectivity with `ping` (see the troubleshooting section of *Installing and Administering LAN/9000 Software*).

1B. List current servers. List the servers currently running on your system by executing the following:

```
netstat -a
```

Table 5-3 lists the servers required for each service.

Table 5-3 Servers Required for Each Service

Local Address	Client/Request	TCP State
*.ftp	ftp	LISTEN
*.telnet	telnet	LISTEN
*.login	rlogin	LISTEN
*.shell	remsh, rcp	LISTEN
*.exec	rexec library	LISTEN
*.who	rwho, ruptime	
*.smtp	sendmail SMTP	LISTEN
*.tftp	tftp	LISTEN
*.bootps	bootpd	LISTEN
*.finger	fingerd	LISTEN

UDP-based protocols are datagram driven, so they do not show a TCP LISTEN status.

1C. Server exists for service? If the server does not exist for the requested service, continue with 1D to determine why. If the server does exist for the server, continue with 1C1.

1C1. Go to Flowchart 2. Go to the next flowchart to begin troubleshooting the security of the Internet Services.

1D. Are files correct? Is there an entry for the servers or services in the /etc/inetd.conf or /etc/services files?

Table 5-4 lists the entries that are required in the /etc/inetd.conf file.

Table 5-4 Entries Required in /etc/inetd.conf

Service Requested	inetd.conf Entry
ftp	ftp stream tcp nowait root /usr/lbin/ftpd ftpd
telnet	telnet stream tcp nowait root /usr/lbin/telnetd telnetd
rlogin	login stream tcp nowait root /usr/lbin/rlogind rlogind
remsh, rcp	shell stream tcp nowait root /usr/lbin/remshd remshd
rexeclibrary	exec stream tcp nowait root /usr/lbin/rexecd rexecd
tftp	tftp dgram udp nowait root /usr/lbin/tftpd tftpd
bootpd	bootps dgram udp wait root /usr/lbin/bootpd bootpd
fingerd	finger stream tcp nowait bin /usr/lbin/fingerd fingerd

Check the permissions on the files in the /usr/lbin and /usr/sbin directories. The files ftpd, bootpd, telnetd, rlogind, remshd, rexecd, rwhod, and inetd must be owned and executable by root only. The file fingerd must be owned and executed by bin only. No other user should have permission to write them, although all users can read them.

Table 5-5 lists the entries that are required in the /etc/services file.

Table 5-5 Entries Required in /etc/services

Service Requested	/etc/services Entry
ftp	ftp 21/tcp
telnet	telnet 23/tcp
sendmail/SMTP	smtp 25/tcp
rexec library	exec 512/tcp
rlogin	login 513/tcp
remsh and rcp	shell 514/tcp
rwho and ruptime	who 513/tcp
tftp	tftp 69/udp
bootpd	bootps 67/udp and bootpc 68/udp
fingerd	finger 79/tcp

If the file entries or permissions are not correct, continue with 1E.

1D1. Issue the ps command to check for the Internet daemon. To see if the inetd daemon is active on the server node, log on to the server node and execute the following:

```
ps -ef | grep inetd
```

1D2. The ps command lists only the grep process for inetd? If the grep message is the only response, inetd is not active. If this is true, continue with 1D3.

1D3. Start the Internet daemon. To start inetd, execute the following as superuser:

```
/usr/sbin/inetd
```

Alternatively, if you want to start connection logging, run the following command:

```
/usr/sbin/inetd -l
```

The /sbin/init.d/inetd shell script usually starts inetd at boot time.

1D4. Go to 1B. After `inetd` is running, repeat this flowchart beginning with 1B.

1E. Correct the files. If there was an incorrect entry or no entry in the `/etc/inetd.conf` or `/etc/services` files, enter the correct information and continue with 1D1.

1F. Reconfigure the Internet daemon. To reconfigure `inetd`, execute the following as superuser:

```
/usr/sbin/inetd -c
```

Continue with 1G.

1G. Go to 1B. Repeat flowchart from 1B to check if the server exists.

Flowchart 2. Security for telnet and ftp

Even though a server exists for a service, the server may not accept connections due to the security that has been implemented for the server.

Follow Flowchart 2 to troubleshoot security for telnet and ftp services.

Figure 5-4 Flowchart 2. Security for telnet and ftp

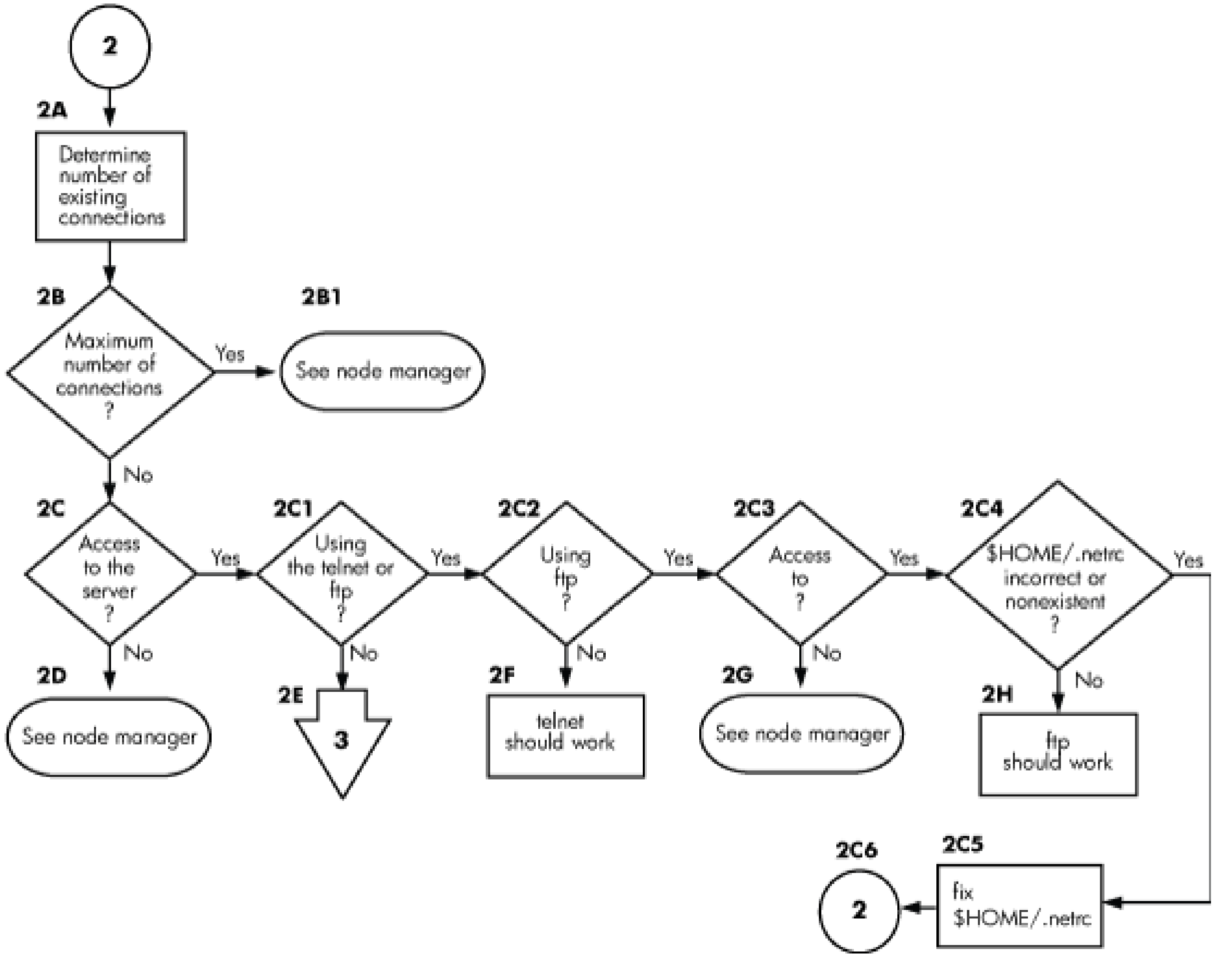

NOTE: The corrections suggested in 2B1, 2C1, and 2F1 must be done by the superuser. Also, except for the anonymous user ID, `ftp` requires non-null passwords on remote user accounts.

2A. Determine the number of existing connections. If `inetd` was started with the `-l` option, the system log may list the number of connections. If these messages do not appear in the system log, continue with 2B, or enable the connection logging with `inetd -l`.

2B. Maximum number of connections? The maximum number of simultaneous connections is specified in the optional file `/var/adm/inetd.sec`. When `inetd` is configured, it checks this file to determine the number of allowable incoming connections. Look at this file to determine how many connections are allowed. The default is 1000.

2B1. See the node manager. If the maximum number of connections has been reached, the node manager can change this value in the `/var/adm/inetd.sec` file.

2C. Access to the server? The `/var/adm/inetd.sec` file also contains a list of systems that may not access the server. If `inetd` was started with the `-l` option, the system log may list the connections that are refused access to the server. Check this log file, if it exists, or ask the node manager to verify whether you have access to the server. If you find that you do not have access to the server, continue with 2D.

2C1. Using `telnet` or `ftp`? Additional security files exist for these services that must be checked. If you are using `ftp` or `telnet` go to 2C2; otherwise, go to 2E.

2C2. Using `ftp`? If you are attempting to use `ftp`, go to 2C3; otherwise, go to 2F.

2C3. Access to `ftp`? If the user you are logging on as is listed in the `/etc/ftpusers` file on the server system, you may not use `ftp` to that system. If you do not have access to `ftp`, go to 2G.

2C4. `$HOME/.netrc` file incorrect or non-existent? If this file is incorrect or non-existent, it is not used for the connection attempt. In particular, if the file exists, check its mode bits, owner ID, and syntax. Type `man 4 netrc` for more information. If it is correct, go to 2H.

2C5. Fix `$HOME/.netrc`. If the file is incorrect, make corrections to it and go to 2C6.

2C6. After you have made the corrections, repeat this flowchart beginning with 2A.

2D. See the node manager. If your system was denied access to the server system by the `/var/adm/inetd.sec` file, but you want to use the server, contact the node manager of the server system and request access.

2E. Go to Flowchart 3. If you are using the Berkeley Services (`sendmail`, BIND, `finger`, the `rexec` library, or any of the “`r`” services), go to Flowchart 3 to begin troubleshooting the security for those services.

2F. `telnet` should work. If you have reached this point in the flowchart, the `telnet` server exists and you have access to the system. If you are using correct syntax, if the login password you are using exists on the server system, and if none of the error messages have solved the problem, report the problem to your Hewlett-Packard support contact.

2G. See the node manager. You are not allowed to use `ftp` to access the server system. Check with the node manager of the server system and request that the appropriate user name be removed from the `/etc/ftupusers` file.

2H. `ftp` should work. If you have reached this point in the flowchart, the `ftp` server exists and you have access to the system. If you are using correct syntax and none of the error messages have solved the problem, report the problem to your Hewlett-Packard support contact.

Flowchart 3. Security for Berkeley Services

Flowchart 3 is for troubleshooting security for the Berkeley Services: `sendmail`, BIND, `finger`, the `rexec` library, and those services that begin with `r`. The following information assumes an account has a password. If it does not, the security checks are not performed.

Figure 5-5 Flowchart 3. Security for Berkeley Services

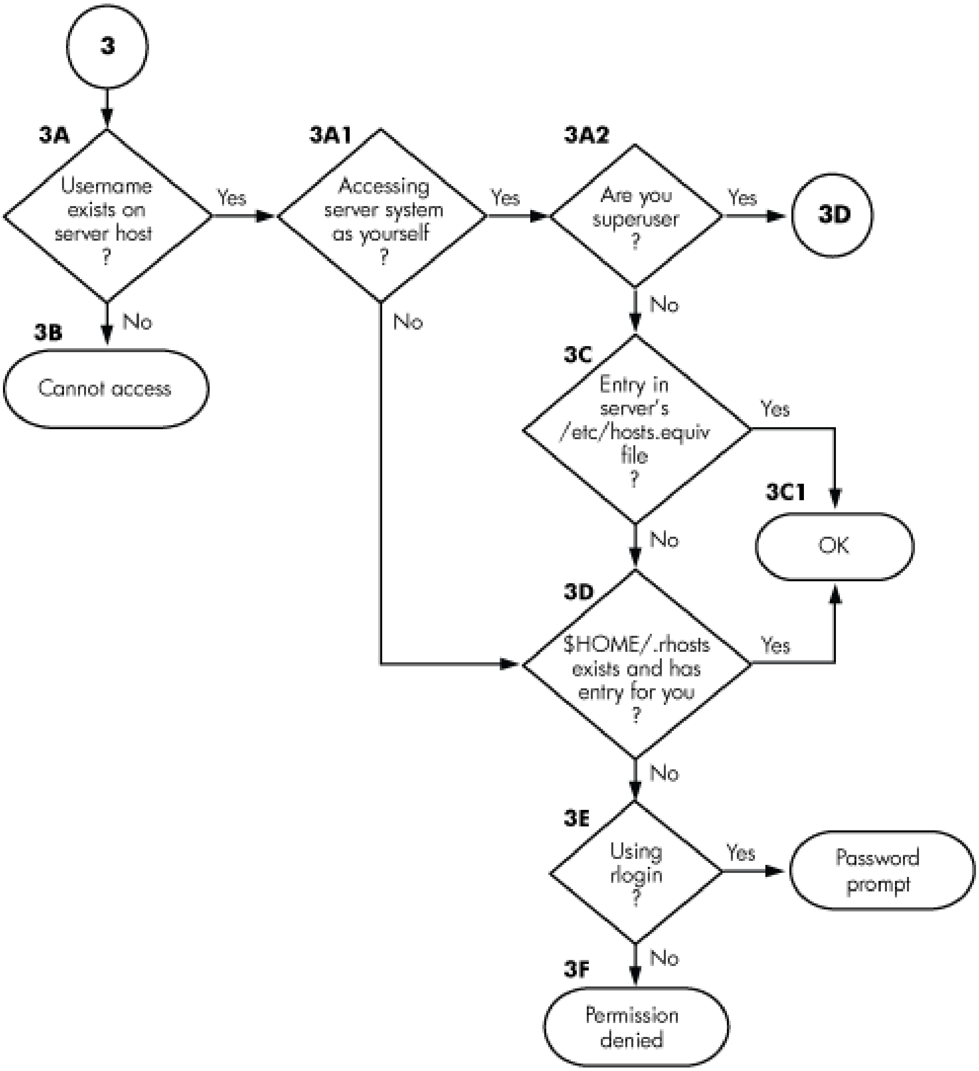

3
3A
Username exists on server host ?
Yes
No
3B
Cannot access
3A1
Accessing server system as yourself ?
Yes
No
3A2
Are you superuser ?
Yes
3D
No
3C
Entry in server's /etc/hosts.equiv file ?
Yes
3C1
OK
No
3D
$HOME/.rhosts exists and has entry for you ?
Yes
No
3E
Using rlogin ?
Yes
Password prompt
No
3F
Permission denied

3A. User name exists on server host? Does the user name that you want to log in as exist on the server host? You can specify another user's name by using the `-l` option with `rlogin`. If the desired user name does not exist on the server host, continue with 3B.

3A1. Accessing server system as yourself? If not, go to 3D.

3A2. Are you superuser? If you are, go to 3D; otherwise, continue with 3C.

3B. Cannot access. Because your user name or the user name that you want to use to log on does not exist on the remote system, you cannot log on to the remote system unless the remote system's node manager creates an account for you.

3C. Entry in server's `/etc/hosts.equiv` file? Does the server system have your official host name entered in its `/etc/hosts.equiv` file? If so, you should be logged on to the remote system without a password prompt. If you can do this, continue with 3C1; otherwise, go to 3D.

3C1. OK. If you are using the `rlogin` service, you are automatically logged in. If you are using another Berkeley service, permission is granted for the operation.

3D. `$HOME/.rhosts` file exists and has entry for you? Does the user name that you want to become on the server system have a `.rhosts` file in that user's `$HOME` directory? If it does, does it list your local host and user name properly? If the `$HOME/.rhosts` file does not exist on the server system, or if it does not have an entry for you, continue with 3E; otherwise, continue with 3C1.

3E. Using `rlogin`? If you are using the `rlogin` service go to 3E1. If you are not using `rlogin`, go to 3F.

3E1. Password prompt. You will receive a password prompt. Enter the password for your remote user name.

3F. Permission denied. You do not have permission to access the user's account. Ask the user to add your local host and user name to his or her `.rhosts` file.

NOTE: For C2 Security, see *A Beginner's Guide to HP-UX* and the *HP-UX System Security Manual*.

Reporting Problems to Your Hewlett-Packard Support Contact

If you do not have a service contract with HP, you may follow the procedure described in this section, but you will be billed accordingly for time and materials.

If you have a service contract with HP, document the problem as a Service Request (SR) and forward it to your Hewlett-Packard support contact. Include the following information where applicable:

- A characterization of the problem. Describe the events leading up to and including the problem. Attempt to describe the source of the problem. Describe the symptoms of the problem and what led up to the problem.

  Include the following in your characterization: HP-UX commands, communication subsystem commands, job streams, result codes and messages, and data that can reproduce the problem.

Illustrate as clearly as possible the context of any messages. Prepare copies of information displayed at the system console and user terminal.

- Obtain the version, update, and fix information for all software.

  To check your Internet Services version, execute the `what` *`service_name`* command, where *`service_name`* is a network service specific to the networking product, such as `ftp` for Internet Services.

  To check the version of your kernel, execute `uname -r`.

  This allows your support contact to determine if the problem is already known, and if the correct software is installed at your site.

- Record all error messages and numbers that appear at the user terminal and the system console.
- Save all network log files.

  Prepare the formatted output and a copy of the log file for your Hewlett-Packard support contact to further analyze.

- Prepare a listing of the HP-UX I/O configuration you are using for your Hewlett-Packard support contact to further analyze.
- Try to determine the general area within the software where you think the problem exists. Refer to the appropriate reference manual and follow the guidelines on gathering information for problems.
- Document your interim or “workaround” solution. The cause of the problem can sometimes be found by comparing the circumstances in which it occurs with the circumstances in which it does not occur.
- Create copies of any Internet Services or other trace files that were active when the problem occurred for your Hewlett-Packard support contact to further analyze.
- In the event of a system failure, a full memory dump must be taken. Use the HP-UX utility `/sbin/savecore` to save a core dump. See the *HP-UX System Administration Tasks* manual for details. Send the output to your Hewlett-Packard support contact.

Index