HEWLETT PACKARD

HP AdvanceNet

Installing and Administering NFS Services

HP Computer Museum
www.hpmuseum.net

For research and education purposes only.

HP 9000 Computers

Installing and Administering NFS Services

Customer Order Number: B1013-90009
Printed in U.S.A., February 1991

Notice

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Restricted Rights Legend

Use, duplication or disclosure by the Government is subject to restrictions as set forth in paragraph (b)(3)(B) of the Rights in Technical Data and Software clause in DAR 7-104.9(a).

This software and documentation is based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California.

DEC® and VAX® are registered trademarks of Digital Equipment Corp.

NOTICE TO USERS

The Network Information Service (NIS) was formerly known as Yellow Pages (YP). The functionality of the two remains the same, only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc.

Hewlett-Packard Company
19420 Homestead Road
Cupertino, CA 95014 U.S.A.

Printing History

First Edition February 1991

Table of Contents

Documentation Overview

Before reading this manual, you should be familiar with HP-UX and have access to *HP-UX Reference* manuals.

Note The information contained in this manual applies to the HP 9000 Series 300, 400, 600, 700, and 800 computers. Any differences in the installation, configuration, operation, or troubleshooting of these computers are specifically noted.

Except for the "NIS Configuration and Maintenance" chapter, all references to servers and clients apply to NFS servers and clients unless otherwise specified.

You will find this manual helpful if you have any of the following responsibilities for the NFS (Network File System) Services product:

- Installation.
- Initial configuration of NFS, NIS (Network Information Service), VHE (Virtual Home Environment), and REX (Remote Execution Facility) services.
- Routine administration and maintenance of NFS, NIS, VHE, or REX.
- Troubleshooting common NFS, NIS, VHE, or REX problems.

Note If you are using NFS Services, but have no administrative responsibilities, refer to the *Using NFS Services* manual.

Contents of This Manual

Refer to the following list for a brief description of the information contained in each chapter and appendix.

Chapter 1: Documentation Overview

This chapter describes who should use this manual, what is in this manual, and where to go for more information.

Chapter 2: NFS Services Overview

This chapter provides a brief overview of the NFS Services product, particularly the NFS, RPC, RPCGEN, REX, Network Lock Manager, NIS, and VHE services. It also describes common terms and concepts.

Chapter 3: Installation

This chapter explains how to install the NFS Services product.

Chapter 4: NFS Configuration and Maintenance

The first section explains how to set up your files in the correct configuration. It also describes NFS daemons, servers, and file systems.

The second section explains procedures for maintaining an efficient system. It includes topics such as NFS file access removal and clock skew problems.

Chapter 5: Remote Execution Facility (REX)

This chapter explains how to configure and use the Remote Execution Facility (REX). You can use REX to execute commands on a remote host.

Chapter 6: Network Lock Manager

The Network Lock Manager and the Status Monitor permit cooperating processes to synchronize access to shared files via System V file locking primitives. This chapter describes the Lock Manager in detail.

Chapter 7: NIS Configuration and Maintenance

The first section explains how to set up your files in a configuration that allows you to centrally administer your NIS databases.

The second section explains procedures for administering and maintaining NIS. It includes topics such as modifying your system to use NIS and changing your NIS password.

Chapter 8: VHE Configuration and Maintenance

This chapter explains how to configure your system to use the Virtual Home Environment (VHE) service. VHE allows you to set up remote login environments to resemble home node login environments.

Chapter 9: Troubleshooting

This chapter describes how to locate and eliminate network problems, specifically those related to the NFS, NIS, VHE, and REX services.

Appendix A: HP NFS Services vs. Local HP-UX

This appendix describes the basic differences between NFS Services and local HP-UX operations.

Appendix B: Moving From RFA to NFS

This appendix describes how to translate RFA applications to NFS applications.

Appendix C: NFS in an HP-UX Cluster Environment

This appendix lists the interactions between NFS Services and HP-UX cluster nodes.

Appendix D: Password Security

This appendix explains the use of encrypted passwords and password security.

Appendix E: Relinking Applications with RPC Library Functions

This appendix describes the changes made to the RPC library functions at HP-UX Release 8.0 and the impact these changes have on applications.

Glossary

The glossary lists and defines terms used in this manual that may not be familiar to you.

Index

The index provides a page reference to the subjects contained within this manual.

Conventions

Table 1-1 explains the conventions used in this manual.

Table 1-1. Conventions

Notation	Description
Boldface	Boldface type is used when a term is defined.
`Computer Text`	Computer type is used for commands and keyboard entries that you must type exactly as shown. It is also used for on-screen prompts and messages.
italics	Italic type is used for emphasis and for titles of manuals and publications. Italic type is also used to represent a variable, such as *user_login_name*.
[Key]	This font is used to indicate a key on the computer's keyboard. When two or more keys appear together with dashes separating them, such as [Ctrl]-[D], press those keys simultaneously to execute the command.
Softkey	This font is used to represent function softkeys that appear at the bottom of your screen.
Underlining	Underlining is used to emphasize a user entry. It distinguishes what you type, such as a command, from other data on the command line, such as the command prompt, a computer response, or a variable. For example: $ date
[ ]	An element inside brackets in a syntax statement is optional.
. . .	A horizontal ellipsis in a syntax statement indicates that a previous element can be repeated. For example: [*option*][*option*]...

Documentation Guide

For More Information	Read
ARPA Services: Daily Use	*Using ARPA Services*
ARPA Services: System Administration	*Installing and Administering ARPA Services*
C2 Security	*HP-UX System Security Manual* *HP-UX Beginner's Guide* *A Beginner's Guide to Using Shells*
Commands and System Calls	*HP-UX Reference Manual*
Network Services: Daily Use	*Using Network Services*
Network Services: System Administration	*Installing and Administering NS Services*
Networking: General Information	*Networking Overview*
NFS Services: Common Commands	*Using NFS Services*
NFS Services: Programming and Protocols	*Programming and Protocols for NFS Services*
NFS Services: System Administration - Configuration - Installation - Maintenance - Network Information Service (NIS) - Network Lock Manager - Remote Execution Facility (REX) - Troubleshooting - Virtual Home Environment (VHE)	*Installing and Administering NFS Services*

Military Standards and Request for Comment Documents

To obtain information about available RFCs, contact the:

Network Information Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

To obtain information about available MIL-STD specifications, contact:

Department of the Navy
Naval Publications and Forms Center
5801 Tabor Avenue
Philadelphia, PA 19120-5099

2

NFS Services Overview

HP's NFS (Network File System) Services product allows many systems to share the same files. It is an independent networking product, not a distributed operating system. NFS differs from distributed operating systems by not limiting its use to specific hardware and software. Rather, it operates on heterogeneous nodes and in operating systems from a variety of vendors. Explicit file transfers across the network to your local node are unnecessary. Since access techniques are transparent, remote file access remains similar to local file access.

With NFS all network nodes are either clients or servers or both, as defined below:

- A **client** is any node or process that accesses a network service.

  An NFS client can also be configured as any combination of an NFS server, NIS (Network Information Service) client, or NIS server. (An NIS server must also be configured as an NIS client.)

- A **server** is any node that provides one of the network services. A single node can provide more than one service.

  An NFS server can also be configured as any combination of an NFS client, NIS client, or NIS server. (An NIS server must also be configured as an NIS client.)

  Servers are passive in that they always wait for clients to call them. The degree to which clients **bind** to their server varies with each of the network services. However, the client always initiates the binding. The server completes the binding subject to access control rules specific to each service.

  NFS servers are **stateless**; they do not maintain information relating to each client being served. Each file request goes to the appropriate server with the parameters attached to it locally (e.g., read and write privileges). An advantage of servers being stateless is that you can reboot servers without adverse consequences to the client.

Components of the NFS Services

The NFS Services product includes the following components:

- NFS remote file access.
- Remote Execution Facility (REX).
- Remote Procedure Calls (RPC).
- Remote Procedure Call protocol compiler (RPCGEN).
- External Data Representation (XDR).
- Network Lock Manager.
- Network Status Monitor.
- Network Information Service (NIS).
- Virtual Home Environment (VHE).

The NFS, REX, Lock Manager, and NIS functionalities are built on top of RPC and XDR library routines.

NFS Remote File Access

Before the client can access remote files, the following steps must be done:

- On the server, the superuser must export the file system (i.e., make it available) to the client.
- On the client, the superuser must mount (import) the file system.

Note

Like local HP-UX operations, if you copy files from a long file name file system to a short file name file system, then file names longer than 14 characters will be truncated after the 14th character.

Long and short file name file systems are set up by the System Administrator.

Access to remote files is the same as for local files. You need to include either the complete path name starting with / (slash) or the path name relative to the current directory. The following figure and steps explain how NFS remote file access works.

EXAMPLE:

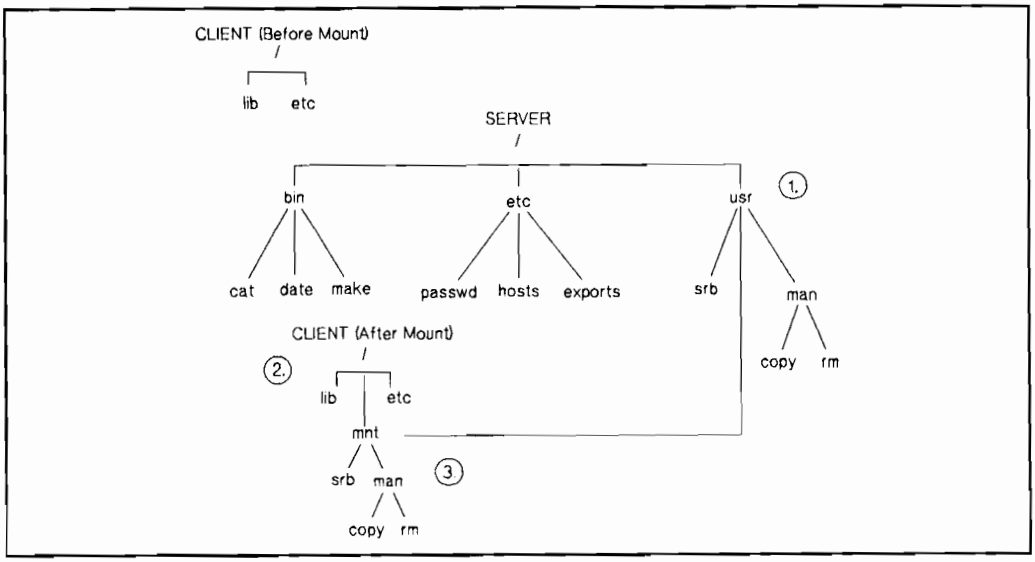

Figure 2-1. NFS Remote File Access

1. The superuser edits the server's `/etc/exports` file to make the `/usr` file system available to the client.

   ```
   server superuser% cat /etc/exports
   /usr client_name
   ```

2. On the client, the superuser creates a mount point `/mnt` (empty directory) and mounts the file system.

   ```
   client superuser% mkdir /mnt
   client superuser% mount server:/usr /mnt
   ```

3. The client reads the files in the `/mnt` directory.

   ```
   client% more /mnt/man/copy
   ```

Two very important features of NFS Remote File Access are **named pipes** and **device files**. The following sections explain the details of these two features.

Named Pipes

A named pipe is a special type of object in the HP-UX file system. A named pipe is one of the many ways in HP-UX that unrelated processes can communicate. HP-UX processes executing on the same client system are able to communicate using named pipes. You can use named pipes via normal file operations, e.g. `open()`, `close()`, `read()`, `write()`. Typically, one process will open the named pipe for reading and another process will open it for writing.

To illustrate named pipes, consider the following example:

EXAMPLE:

> C1 and C2 are processes executing on system C. Also assume host C has mounted file system / from host S on `/mnt`. C1 opens `/mnt/FIFO` for reading and C2 opens `/mnt/FIFO` for writing. C1 can now read what C2 wrote to the named pipe.
>
> Next, assume a third process (process D3) is running on another client D which also has / from S mounted on `/mnt` (on system D), and it opened `/mnt/FIFO` for reading. Is process D3 able to read what process C2 wrote to this named pipe? No, because no actual NFS activity occurs between the NFS client and NFS server for named pipe reads and writes. These are handled entirely by the client.

Note In certain cases there would be NFS activity. For example, if you do a `chown` on the named pipe, the request will go to the server to change the owner.

---

mknod()

Named pipes are created with `mknod()`. Any user can create a named pipe with `mknod()`. (Use of `mknod()` to create device files requires superuser privileges.)

Note If you attempt to make a directory or a network special file over NFS, `mknod()` will fail and will return with `errno` set to EINVAL (invalid argument).

Device Files

Device files are another type of object in the file system, and are used to access physical or conceptual devices attached to the system. NFS device files always refer to a device attached to the local system and can generally be used where a local device file would be used. Like named pipes, device files are operated on through normal file system operations. For example, to write to the system console, you can write to the file `/dev/console`.

To illustrate the use of device files, consider the following:

EXAMPLE:

> System C is an NFS client of the NFS server system S, and has mounted file system / from host S on `/mnt` (a superuser on system C executed the command `mount S: / /mnt`). If a process on system C attempts to write to `/mnt/dev/console`, a device file representing the system console on system S, the output will go to the system console on system C, not on system S. If a process on system S attempts to write to `/dev/console`, which is the same "file" that system C wrote to, it will actually write to the console on system S.

NFS Mounts: Turning Off Device File Access

NFS device files are not secure. Therefore, the system administrator has the option of turning off device file access on a per-NFS mount basis. The administrator uses the `-o nodevs` option to the `mount` command to turn off device file access.

EXAMPLE:

```
mount -o nodevs nfserver:/servermountpoint /clientmountpoint
```

Note The nodevs option does not turn off support of named pipes.

NFS Mounts: Mounting From NFS Device Files

You may mount a local disk that is represented by a remote NFS device file.

EXAMPLE:

```
mount /mnt/nfs/dev/dsk/0s0 /localmntpt
```

Access to the newly mounted file system will proceed as if the disk had been mounted from a local device file.

Note Access to the local disk's mounted file system will not be affected even if the NFS file system is unmounted.

Normally when unmounting a file system, you can give either the name of the device file or the name of the mount point. However, if the NFS server is down or the NFS file system is down, you must give the mount point to unmount the local disk.

EXAMPLE: You would enter the following to unmount a local disk:

```
umount /localmntpt
```

instead of:

```
umount /mnt/nfs/dev/dsk/0s0
```

The latter case will not fail if the NFS server is down, but it will hang until the server comes back up as any other NFS access does.

Remote Execution Facility (REX)

The Remote Execution Facility allows you to execute commands on a remote host. REX is similar to the Berkeley Service remote shell (`remsh`) with two major differences:

- Your environment is simulated on the remote host.
- You can execute interactive commands on the remote host.

Remote Procedure Call (RPC)

NFS Services consists of remote programs composed of remote procedures called from the client nodes on the network. Optimally, a remote procedure computes results based entirely on its own parameters. Thus, the procedure (and therefore, the network service) is not tied to any particular operating system or hardware.

NFS clients access server information and processes by making a remote procedure call. RPC allows a client process to execute functions on a server via a server process. Though these processes can reside on different network nodes, the client process does not need to know about the networking implementations.

The client first calls an RPC function to initiate the RPC transaction. The client system then sends an encoded message to the server. This message includes all the data needed to identify the service and user authentication information. If the message is valid (i.e., calls an existing service and the authentication passes) the server performs the requested service and sends a result message back to the client.

Remote Procedure Call Protocol Compiler (RPCGEN)

RPCGEN is a Remote Procedure Call compiler. You use it to convert applications running on a single computer to ones that run over a network. It is also used to assist in writing Remote Procedure Call applications simply and directly. With RPCGEN, your development time will be reduced and you will spend less time coding and debugging network interface code.

You produce three of the files required to convert an application to run on a network. These files are:

- Protocol description file.
- Client side file.
- Server side function file.

RPCGEN accepts remote program interface definitions (the protocol description file) written in RPC and produces the following C output files, which you may use as a starting point, rewriting as necessary:

- Header file.
- Client side subroutine file.
- Server side skeleton file.
- XDR (External Data Representation) routine file.

If you wish to use the RPCGEN compiler to write RPC applications, refer to the "RPCGEN Programming Guide" chapter in the *Programming and Protocols for NFS Services* manual.

External Data Representation (XDR)

RPC uses the eXternal Data Representation functionality to translate machine dependent data formats (i.e., internal representations) to a universal format used by all network nodes using RPC/XDR. Thus, XDR enables heterogeneous nodes and operating systems to communicate with each other over the network.

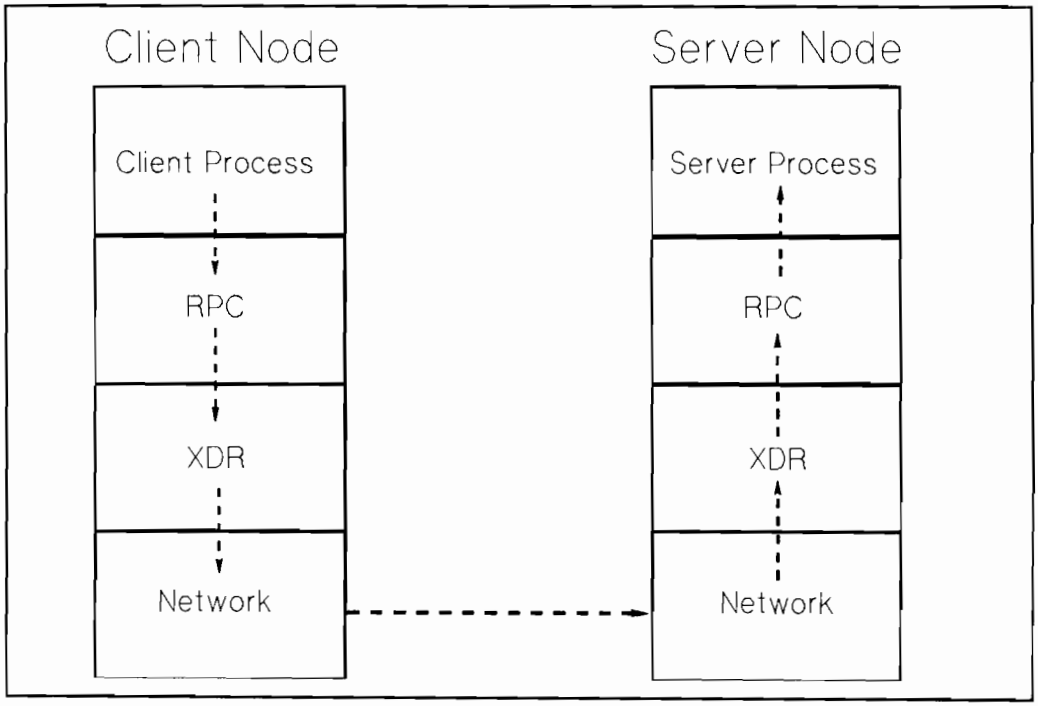

Figure 2-2. RPC and XDR Data Transfer

Note This figure does not correspond to the ISO Model.

Network Lock Manager and Network Status Monitor

NFS Services includes the Network Lock Manager (`rpc.lockd`) and the Network Status Monitor (`rpc.statd`). The Network Lock Manager supports file locking and synchronized access to shared files via `lockf` and `fcntl` for NFS. The Network Status Monitor is used by the Network Lock Manager to maintain the stateful locking service within the stateless NFS environment. It allows applications to monitor the status of other computers and systems.

Network Information Service (NIS)

The Network Information Service (NIS) is an optional service containing a collection of cooperating NIS server processes that provide NIS clients access to data. (NIS was formerly known as Yellow Pages (YP) which is a registered trademark of British Telecommunications.) You can administer all the databases from one NIS *master server* since it propagates data across the network to other NIS servers. NIS includes the following features:

- NIS manages unlimited databases. Typically these include files in `/etc/group`, `/etc/hosts`, `/etc/netgroup`, `/etc/networks`, `/etc/passwd`, `/etc/protocols`, `/etc/rpc`, and `/etc/services`.

  For example, programs previously read `/etc/hosts` to find an Internet address that corresponds to a host name. When you added a new node to the network, you had to add a new entry to every node's `/etc/hosts` file. Now programs can use NIS to obtain information from other NIS servers.

- Since the NIS master server propagates all *maps* (databases) to the **slave servers**, an NIS client receives consistent information regardless of which NIS server it accesses.

- If a remote node running an NIS server process crashes, NIS client processes can obtain NIS services from another NIS server.

- Since the NIS interface uses RPC and XDR, the service is available to other vendors.

NIS Advantages

NIS has several advantages:

- NIS enables you to automatically keep user IDs and group IDs consistent among all the nodes participating in NFS file sharing.

  Without NIS, you have to manually keep these IDs consistent for NFS.

- NIS provides the convenience of centrally administering the `/etc` files: `group`, `hosts`, `netgroup`, `networks`, `passwd`, `protocols`, `rpc`, and `services`.

  Without NIS, you must administer these files on each node individually.

NIS Disadvantages

NIS has the following disadvantages:

- If a network grows beyond 2000 nodes, NIS may begin to exhibit poor performance or failures. (This limit is based on todays system capacity.)
- Since NIS provides NIS clients access to data via the network, NIS clients may observe slower performance than if the data were accessed from local files. For example, with NIS, logging in may take more time if the NIS server is busy.
- If any of the NIS servers are unstable, remote access to files may be slower since the NIS client may have to rebind to another NIS server. If no other NIS server is available, users may not be able to login to their nodes without access to the NIS's `passwd` map.
- NIS does not make changes visible to all users unless the changes are made on the NIS master server.
- The NIS slave servers do not immediately see the changes made to the NIS master server maps. The updated maps become consistent among all NIS servers only after each slave server successfully copies the maps via `ypxfr`.

Note If you configure the BIND Name Server, it will be used instead of NIS for host name and address resolution. NIS will still be used for all other information such as passwords. See "Configuring and Maintaining the BIND Name Server" in the *Installing and Administering ARPA Services* manual.

NIS Concepts

Refer to the following figure and subsections for a summary of how components within the Network Information Service work together: maps, NIS domains, NIS servers (masters and slaves), and NIS clients.

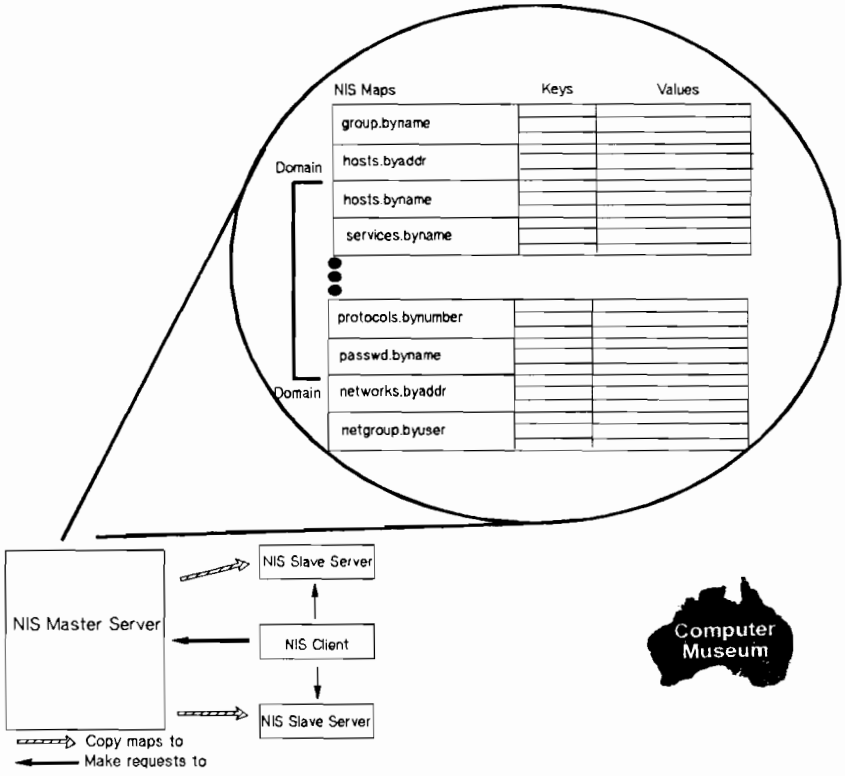

Figure 2-3. Network Information Service Structure

NIS Maps

The NIS system stores information in NIS maps (databases). Each map contains a set of keys and associated values: one key per value and one value per key. (A value may be a string of characters with imbedded blanks or tabs). For example, in the `passwd.byname` map, all the login names are the keys and their matching lines from `/etc/passwd` are the values.

Each map has a unique **map name** that programs use to access the map. Programs must know the format of the data in the map. Many of the maps are derived from ASCII files such as `/etc/hosts`, `/etc/group`, and `/etc/passwd`. The map format is usually identical to the ASCII file format.

NIS Servers and NIS Clients

NIS servers are nodes that provide access to NIS maps via the network. These maps are in `/usr/etc/yp` subdirectories named after the appropriate NIS domains. (See the next section, "NIS Domains.")

NIS clients are nodes that request access to NIS maps from an NIS server as follows:

1. An NIS client that is not bound sends a broadcast to all NIS servers on the network.
2. The NIS client binds to the first NIS server that responds. (Each NIS client binds to one NIS server per NIS domain.)
3. If the request is the NIS client's first attempt to access data, the NIS client remembers which NIS server responded to the request. Subsequent requests by this NIS client go directly to this NIS server.
4. If the bound NIS server is down or unavailable, the NIS client automatically rebinds to the first NIS server that responds to another broadcast.

Note An NIS client can also be configured as any combination of an NIS server, NFS client, or NFS server.

An NIS server must also be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.

NIS Domains

An **NIS domain** is a logical grouping of the set of maps contained on NIS servers. The following rules apply to NIS domains:

- Nodes that belong to the same NIS domain have the same domain name.
- An NIS domain has only one master server.
- An NIS domain may have zero or more slave servers.
- Maps with the same name in different NIS domains can have different contents.

You implement an NIS domain as a subdirectory of `/usr/etc/yp` on each NIS server; the name of this subdirectory is the name of the NIS domain. For example, maps in the `research` NIS domain would be in `/usr/etc/yp/research`. (Note that NIS domain names are case sensitive.) All directories that appear under `/usr/etc/yp` are assumed to be domains that an NIS server serves. To remove a domain being served, you must delete that domain's subdirectory name from `/usr/etc/yp` on all of its servers.

The `/etc/netnfsrc` file usually contains the default NIS domain name. You can change the default by executing the `domainname` command or by editing `/etc/netnfsrc` and then rebooting the system.

NIS Master and NIS Slave Servers

Only two types of nodes have NIS databases: master and slave servers.

The **NIS master server** is the node on which NIS maps are built from ASCII files; it, therefore, contains the master databases (maps) which other NIS servers (slaves) copy. Note that the NIS master server may also provide NIS clients access to NIS maps.

Note You should create and modify NIS databases only on the NIS master server; otherwise, all NIS databases will not be consistent across the NIS servers.

The **NIS slave servers** are the nodes that receive the propagated maps from the NIS master server. In turn, they provide NIS clients access to NIS maps.

An NIS server can be the master or slave of many domains. However, an NIS server can only be either the master or a slave of a given domain.

Though an NIS server may be master for one map and slave for another, random assignment of maps to NIS master servers may cause confusion. Therefore, only one NIS server should be the master for all maps within an NIS domain.

Virtual Home Environment (VHE)

Virtual Home Environment (VHE) is an HP-developed service that allows you to configure your login environment on remote nodes to mirror the login environment on your home node. (Home node refers to the node on which your home directory physically resides.) VHE is an optional service that is available to any HP-UX system that has the NFS product. It may also be used with other UNIX systems that support symbolic links and NFS.

If you find that you never need to work from a remote node, you may want to skip this section.

VHE Advantages

VHE's major advantage is that you can sit down at any remote node (assuming you have login permission), login, and enter into the work environment that is associated with the login on your home node (your home directory as specified in `/etc/passwd`). This includes:

- Home shell configuration (i.e., whichever shell you are configured to use on your home node appears when you login to a remote node).
- Access to files on the file systems exported for VHE on any computers connected with VHE on the network to which you have a login and file access permission.
- Use of previously defined aliases (only for C or K shells) and shell variables.
- Use of customized shell scripts (assuming shells operate similarly on your home node and the node you are currently using).
- Use of compiled files under your home directory from your home node (assuming your home node and the node you are logged into are of the same architecture and operating system).

Thus, VHE allows you to minimize the number of computer interfaces you must learn to be productive on the various computers that are running NFS on your network and you are no longer tied to a particular computer to complete your work tasks.

Another advantage of VHE is that it distributes computational work more efficiently between nodes than ARPA/Berkeley terminal emulation services such as `telnet` or `rlogin`.

Unlike `telnet` or `rlogin`, VHE does not return to your home node, that contains your home environment login, to execute tasks.

Instead, VHE takes advantage of the computing capacity of the machine you are currently using. For example, if you use VHE on a node other than the home node and perform an `ls` command of a directory on the home node, the `ls` command is executed from the local `/bin` directory. VHE does not return to your home node's `/bin` directory to execute the `ls` command. The following figure illustrates this concept.

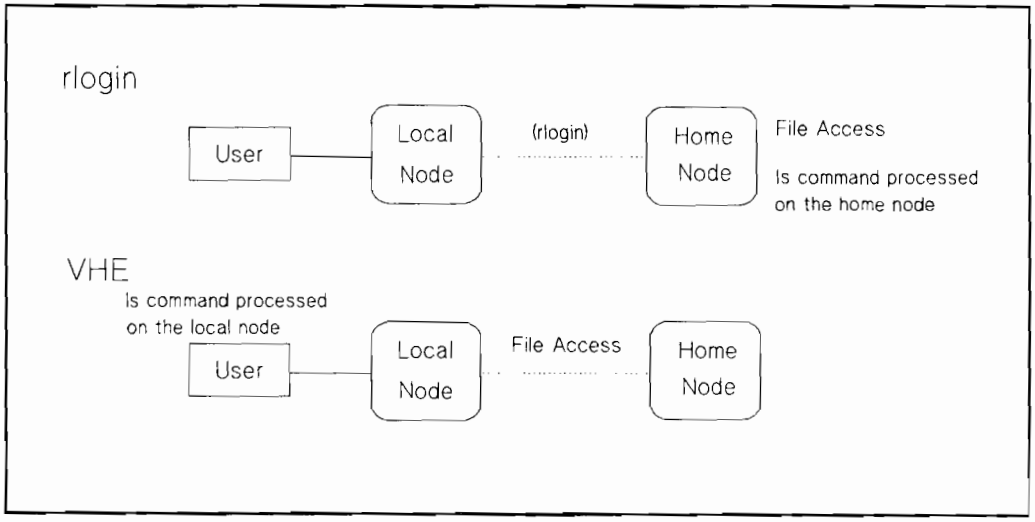

Figure 2-4. VHE vs. rlogin Performing ls Command

VHE Disadvantages

VHE has the following disadvantages:

- Though you can edit source code files originating from different types of computers on the network, you will not be able to execute object code files from a computer of a different architecture using VHE. For example, consider the following: You are currently working on an HP 9000 Series 300 and running VHE, and your home node is an HP 9000 Series 800 computer. If you try to execute an object code file on the HP 9000 Series 300 from the Series 800 computer it will not succeed. However, you can execute a script from the Series 800 computer.

- If you specify pathnames or hardware attributes in your node's `.profile` or `.login` files, you may have to modify these files to use VHE effectively. For example, the `.login` file needs to prompt for the terminal type if you plan to use VHE from more than one terminal or display type. If you do not already have this capability, then look in the sample `/etc/d.login` or `/etc/d.profile` files for samples of how to do this.

- When you are in your home environment, you may execute set-uid root programs that access files in your home directory. These files must allow access for the user "nobody." If this is not done, set-uid root programs will fail. The same applies for root access via set-uid. For example, your home directory is accessed via VHE and you execute set-uid to gain superuser privileges. If your shell happens to be `ksh`, your root `ksh` may hang if your `.history` file does not allow access for user "nobody."

How VHE Works

The following diagram illustrates the directory structure of nodes in a network using VHE.

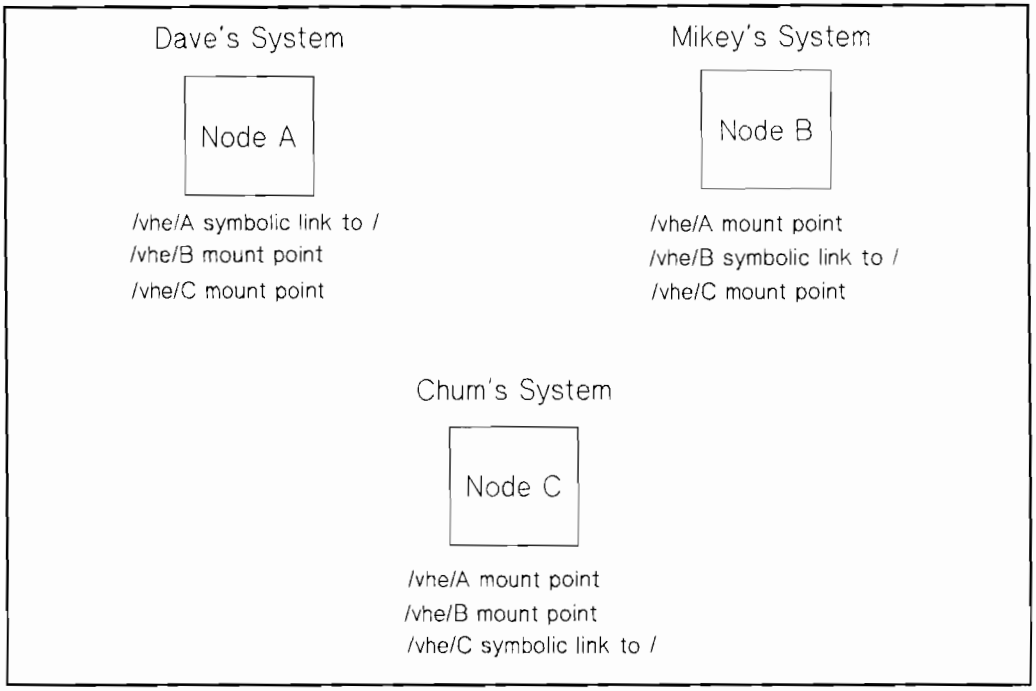

Figure 2-5. Directory Structures of Nodes Using VHE

Each node is connected to the others via NFS Services. In the picture, each node is a home node for a different user (Dave, Mikey and Chum). Each user has a customized work environment set up by the login process. Directories on each home node correspond to each of the remote nodes. For example, on node A there is a directory `/vhe/B` that corresponds to node B. Using these directories as mount points, a mount is done by each node to each remote node. (The definitions of mounts and mount points are included in the "Glossary." More detailed information is contained in the "NFS Configuration and Maintenance" chapter in the *Installing and Administering NFS Services* manual).

Using VHE gives each node access to file systems located on the remote nodes. To maintain consistency when you log into your home node, a symbolic link (a pointer) points to the host's root directory.

In a single node HP-UX configuration, the `/etc/passwd` file contains the directory that becomes the home directory for the user upon logging in. For use with VHE, `/etc/passwd` is edited such that all of the home directories are prefixed with a mount point or a symbolic link. When the login program performs a `cd` to the user's home directory, the `cd` and subsequent requests are made to the user's home node via NFS Services unless logging in on your home node.

Example Grouping

In the `/etc/passwd` file, the appropriate mount point or symbolic link is added to the beginning of the pathname of the home directory for each user. The example below shows how the lines in `/etc/passwd` would look for the users Dave, Mikey, and Chum as shown in Figure 2-5:

```
dave::117:100:Dave:/vhe/A/users/dave:/bin/csh
mikey::118:100:mikey Pom :/vhe/B/users/mikey:/bin/sh
chum::119:200:chum Pom:/vhe/C/users/chum:/bin/ksh
```

No matter which node Dave logs in on, his home directory is `/users/dave` on node A. When scripts such as `.login` or `.cshrc` are executed, they define the execution environment as customized by Dave. His files, shell variables and aliases are available just as if he had physically logged in on node A.

Because VHE is not a virtual terminal program, when Dave executes processes, they are executed on the node he is logged into. If he is on node B, processes are executed on node B, not his native host A. For example, consider the following. Dave is working at node B and his system administrator has configured VHE to be running. Dave enters the following command on node B:

```
cc testfile.c
```

The `cc` from node B's `/bin` directory is executed, but `testfile.c` is used from Dave's current working directory on node A.

Installation

The installation procedures for the HP 9000 Series 300/400 and Series 600/700/800 computers are slightly different. These differences will be noted in the sections that follow.

NFS Installation Checklist

The following steps are a checklist of NFS installation procedures. You may have already completed several of these steps. You will most likely start with Step 4. Steps 4 through 6 are explained in detail in this chapter.

1. Prepare your HP 9000 system for operation:
   a. Inspect hardware.
   b. Create and maintain a network map.
2. Ensure your computer is running the LAN/9000 software. Refer to the *Installing and Administering LAN/9000* manual.

   Ensure that your computer's HP-UX operating system, your LAN/9000 software, and the NFS software that you are about to install all have the same version number. If you do not know which version of HP-UX your computer is running, execute the `uname -r` command.

   If the versions do not match, run `update` to install the correct HP-UX operating system version. Refer to the *System Administration Tasks* manual for information on the `update` procedure.

3. Install the NFS software. You will need to use the `update` program to install the NFS software. Refer to the *System Administration Tasks* manual for detailed `update` information. Do the following:

   a. Use the `/etc/update` command.

   b. If necessary, configure the new kernel to include NFS.

4. Add your HP 9000 computer to the network using your LAN/9000 software. Refer to *Installing and Administering LAN/9000* manual. Do the following:

   a. Assign an internet address.

   b. Edit `/etc/rc` and `/etc/netlinkrc` manually or use SAM (System Administration Manager).

   c. Verify that device files exist for the node's LAN; if they do not, you must create them.

Note If you have applications that access RPC C library routines, it may be necessary to relink those applications. For more information, see Appendix E.

Key Terms

Term	Definition
Cluster	One or more workstations linked together with a local area network (LAN), and sharing a global file system attached to the root server. For more information on cluster concepts, see *Managing Clusters of HP9000 Computers: Sharing the HP-UX Filing System*.
Cluster Auxiliary Server	A cluster client with a disk drive that contains files shared by the other members of the cluster.
Cluster Node (Cnode)	Any node operating in an HP-UX cluster environment, including cluster clients and cluster servers.
Cluster Client	A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have its root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.
Cluster Root Server	The only node in an HP-UX cluster that has the root file system directly attached to it.
Context Dependent File (CDF)	A hidden directory that contains all the versions of a file needed by the different cnodes.
Heterogeneous Cluster	A diskless cluster with more than one type of computer architecture (e.g., Series 300 and Series 800).
Homogeneous Cluster	A diskless cluster composed of nodes of only one computer architecture (e.g., Series 300 only).

Term	Definition
Internet Address	A four-byte quantity that is distinct from a link-level address and is the network address of a computer node. This address identifies both the specific network and the specific host on the network.
LAN	Local Area Network.
Network Information Service (NIS)	An optional network service composed of databases (maps) and processes that provide NIS clients access to the maps. NIS enables you to administer these databases from one node.
NFS	Network File System.
Node	A computer system that is attached to or is part of a computer network.
update	The HP-UX command that installs or updates software onto the system.

Prepare the HP 9000 System

To prepare your HP 9000 computer for operation on the LAN, you must ensure your LAN hardware is installed correctly.

For LAN hardware installation instructions for your computer, refer to the following documentation:

- *LAN Interface Controller (LANIC) Installation and Reference Manual.*
- *Twisted-Pair MAU Installation Guide.*
- *LAN Cable and Accessories Installation Manual.*
- *Installing and Administering LAN/9000*

Another step in preparing your system is to update your network map with all new installation information (e.g., new computers, cable changes). If you do not have a network map, HP strongly recommends you create one. Refer to *Installing and Administering LAN/9000* manual for guidelines.

Install the NFS Software

Before installing NFS Services software, you should ensure the following items are true:

- Your computer's HP-UX operating system, your LAN/9000 software, and your NFS software all have the same version number. Otherwise, the network may malfunction. To check which version of HP-UX you are currently running, execute the `uname -r` command.
- The LAN/9000 software is installed. To verify whether the LAN/9000 software has been installed, check with your systems administrator. If you are the systems administrator, and you have not already installed the LAN/9000 software, refer to the *Installing and Administering LAN* manual for installation and configuration instructions.

Use update Program

Before installing NFS Services, refer to the *System Administration Tasks* manual to familiarize yourself with the `update` program's menu operations and device file information.

After you are certain the required HP-UX and LAN/9000 software is installed, use the `/etc/update` program to install the NFS Services software. The `/etc/update` program takes you through the installation procedure step by step.

After you finish installing the NFS software, log in as superuser and display the /etc/newconfig directory. The installation added the following files to the /etc/newconfig directory. You will use these files when you configure the NFS Services, Network Information Service, and Virtual Home Environment:

- /etc/newconfig/netgroup
- /etc/newconfig/yp_Makefile
- /etc/newconfig/netnfsrc
- /etc/newconfig/ypmake
- /etc/newconfig/rpc
- /etc/newconfig/netnfsrc2
- /etc/newconfig/ypxfr_1perday
- /etc/newconfig/vhe_mounter
- /etc/newconfig/ypxfr_1perhour
- /etc/newconfig/vhe_u_mnt
- /etc/newconfig/ypxfr_2perday
- /etc/newconfig/vhe_list
- /etc/newconfig/ypinit

The configuration procedures are described later in this manual. For descriptions of these files and other files in /etc/newconfig, refer to /etc/newconfig/README.

Note

If you have just updated previously-existing NFS Services on a Series 300/400 system, then you have completed installing the NFS Services product. If you have just updated a Series 300/400 system to add NFS Services for the first time, you must now configure a new kernel to include NFS. To configure a new kernel, refer to the next section.

On a Series 600/700/800 system, you have completed installing the NFS services product, and you do not need to configure a new kernel unless the update program failed to generate a new kernel. To configure a new kernel, refer to the next section.

Configure a New Kernel

To prepare the NFS Services product for use, you must configure a new HP-UX operating system kernel if you are installing NFS Services for the first time.

If your kernel was constructed from the standard kernel file (`/etc/conf/dfile` on a Series 300/400 or `/etc/conf/gen/S800` on a Series 600/700/800), you can use SAM (System Administration Manager) to configure a new kernel that includes NFS.

If your kernel is based upon a customized kernel file, you must manually configure a new kernel. On a Series 300/400, your kernel file must contain the uncommented entry:

```
nfs
```

On a Series 600/700/800, your kernel file must contain the uncommented entry:

```
include nfs;
```

Note If you are configuring NFS in an HP-UX cluster environment, you must configure NFS into the kernel on all cnodes in the cluster.

See the *System Administration Tasks* manual for instructions on configuring a new kernel.

Add a Computer to the Network

If you have not already done so, refer to the *Installing and Administering LAN* manual for instructions on adding your HP 9000 computer to the network. You will need to perform the following steps:

1. Determine and assign an internet address.
2. Edit `/etc/rc` and `/etc/netlinkrc` manually or use SAM (System Administration Manager).
3. Verify that device files exist for the node's LAN; if they do not, create them.

After rebooting the system, log in as superuser and refer to the "NFS Configuration and Maintenance," "NIS Configuration and Maintenance," and "VHE Configuration and Maintenance" chapters to configure your system with NFS, NIS (if applicable), and VHE. Refer to the "Configuring and Maintaining the BIND Name Server" chapter in the *Installing and Administering ARPA Services* manual to configure the BIND Name Server if applicable.

4

NFS Configuration and Maintenance

This chapter describes basic NFS configuration without the Network Information Service (NIS). The latter portion describes how to administer and maintain the NFS service once you have it configured. For specific NFS information, refer to the following sections:

- Key terms.
- Guidelines.
- NFS configuration.
- NFS maintenance.

Note All references to **servers** and **clients** in this chapter apply to NFS servers and NFS clients unless otherwise specified.

Key Terms

Term	Definition
Alias	A term for referencing alternate networks, hosts, and protocol names.
Client	- A node that requests data or services from other nodes (servers). - A process that requests other processes to perform operations. *Note*: An NFS client can also be configured as any combination of an NFS server, NIS client, or NIS server. (An NIS server *must* also be configured as an NIS client.)
Clock Skew	A difference in clock times between systems.
Cluster	One or more workstations linked together with a local area network (LAN), but consisting of only one root file system. For more information on cluster concepts, see *Managing Clusters of HP9000 Computers: Sharing the HP-UX Filing System*.
Cluster Auxiliary Server	A cluster client with a disk drive that contains files shared by the other members of the cluster.
Cluster Client	A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have its root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.
Cluster Node (Cnode)	Any node operating in an HP-UX cluster environment, including cluster clients and cluster servers.
Cluster Root Server	The only node in an HP-UX cluster that has the root file system directly attached to it.
Daemon	Background programs that are always running, waiting for a request to perform a task.
Export	To make a file system available to remote nodes via NFS.
File System	An entire unit (disk partition) that has a fixed size.

Term	Definition
GID	A value that identifies a group in HP-UX.
Hard Mount	A mount that causes NFS to retry a remote file system request until it succeeds, you interrupt it (default option), or you reboot the system.
Host	A node that has primary functions other than switching data for the network.
Import	To obtain access to a remote file system from an outside source; to mount a remote file system.
Internet Address	A four-byte quantity that is distinct from a link-level address and is the network address of a computer node. This address identifies both the specific network and the specific host on the network.
Interruptable Mount	A mount that allows you to interrupt an NFS request by pressing an interrupt key. (Though the interrupt key is not standardized, common ones include [CTRL]-[C] and [BREAK].)
Locally Mounted File System	A file system that is locally mounted on a cluster client in a diskless cluster.
Mount	To obtain access to a remote or local file system or directory (import).
Mount Point	The name of the directory on which a file system is mounted.
Netgroup	A network-wide group of nodes and users defined in */etc/netgroup*.
Network Information Service (NIS)	An optional network service composed of databases (maps) and processes that provide NIS clients access to the maps. NIS enables you to administer these databases from one node. NIS may or may not be active; check with your system administrator.
NFS	Network File System.
NIS Domain	A logical grouping of NIS maps (databases) stored in one location. NIS domains are specific to NIS and are not associated with other network domains.

Term	Definition
Node	A computer system that is attached to or is part of a computer network.
Server	- A node that provides data or services to other nodes (clients) on the network. - A process that performs operations as requested by other processes. *Note*: An NFS server can also be configured as any combination of an NFS client, NIS client, or NIS server. (An NIS server *must* also be configured as an NIS client.)
Soft Mount	An optional mount that causes access to remote file systems to abort requests after one NFS attempt.
UID	A value that identifies a user in HP-UX.
Unmount	To remove access rights to a file system or disk that was mounted via the *mount* command.
update	The HP-UX command that installs software onto the system.

Guidelines

Refer to the following guidelines for information regarding:

- Network memory.
- Configuration files.
- Daemons.
- Servers.

Network Memory

Network memory is configurable using three parameters: `netmeminit`, `netmemmax` and `netmemthresh`. The default values are generally sufficient for most NFS configurations. However, if you change these parameters, do not set `netmemmax` equal to or less than `netmemthresh`.

To check the memory available on your network, enter the following command:

```
netstat -m
```

For more information about network memory configurations, refer to the *System Administration Tasks* manual.

Configuration Files

The following table lists the files that must be configured (unless otherwise stated) for your system to operate correctly. (Refer to the *HP-UX Reference* for detailed information.)

Configuration File	Description
`/etc/checklist`	Contains a list of file systems that are automatically mounted at boot time.
`/etc/exports`	Contains a list of file systems that clients may import. *Note*: Create this file only on servers.
`/etc/inetd.conf`	Contains information about servers started by `inetd`, including RPC services.
`/etc/netgroup`	Contains a mapping of network group names (netgroups) to a set of node, user, and NIS domain names; both `/etc/exports` and `/etc/passwd` can use the netgroups defined in `/etc/netgroup`. Classifies the nodes for remote mounts. For *ARPA Services*, classifies the users for remote logins and remote shells. You can specify netgroups in `/etc/hosts.equiv` and `$HOME/.rhosts`. Configuring this file is optional.

Configuration File	Description
/etc/netnfsrc	Automatically executed at boot time to start the NFS networking (e.g., starts daemons and servers, defines servers and clients). *Note*: For 8.0, /etc/netnfsrc has changed significantly. A new /etc/netnfsrc file is loaded during the install and update procedure. Lines in which configuration variables are set are propagated from the old version to the new version. The old version is saved in /etc/netnfsrc.OLD. You must copy any customization necessary from the old version. In an HP-UX environment, /etc/netnfsrc is a CDF. Therefore, customization to this file on a cnode in the cluster will only affect that cnode's context.
/etc/netnfsrc2	Perform mount operations for all NFS mount entries found in /etc/checklist. This file is static; it is already correctly configured.
/etc/rpc	Maps the RPC program names to the RPC program numbers and vice versa. This file is static; it is already correctly configured.
/usr/adm/inetd.sec	Checks the internet address of the host requesting a service against the list of hosts allowed to use the service. Specifies how many remote users can simultaneously start remote services in the local system and which remote hosts (or networks) can use the system.

Daemons

The following table lists the networking daemons (background programs) that are always running, waiting for a request to perform a task. The parenthetical comments refer to the *HP-UX Reference* sections where you can go for more information.

Daemon	Description
biod(1M)	Asynchronous block I/O daemons for NFS clients.
inetd(1M)	Internet daemon that listens on service ports. It: - Reads `/etc/inetd.conf` to determine the appropriate server for handling the incoming request. - Listens for and accepts network requests. - Invokes the appropriate server. *Note*: Since `inetd` contacts `portmap` on behalf of the servers it starts, you must start `portmap` before starting `inetd`.
nfsd(1M)	NFS server daemon that responds to client file system requests. When a client program needs to read or write in a remote file system, it sends a request to that system's `nfsd` process. If operating in an HP-UX cluster environment, `nfsd` must be running on any cnode with a local file system that will be exported via NFS. Any `nfsd` daemons running on cnodes without locally mounted file systems are ignored.
pcnfsd(1M)	Daemon that authenticates a PC user's access to files. It takes the user name and password, and then does *one* of the following: - Succeeds (returns a valid UID and GID). - Fails (indicates the name and password are unacceptable). *Note*: Though `pcnfsd` enables PC users to use printer spooling facilities on HP-UX systems, they *must* have the appropriate PC networking software product for it to work.

Daemon	Description
portmap(1M)	Daemon that converts RPC program numbers into port numbers. When `inetd` starts, it tells `portmap`: - Which RPC servers it is listening for. - On which ports it is listening. - The RPC program numbers and versions it serves. When a client makes an RPC call to a given program number, it first contacts `portmap` on the server node to determine the port number where RPC requests should be sent. *Note*: Since `inetd` contacts `portmap` on behalf of the servers it starts, you *must* start `portmap` before starting `inetd`.

Servers

The following table lists the networking servers (processes that perform operations as requested by other processes). The parenthetical comments refer to the *HP-UX Reference* sections where you can go for more information.

Server	Description
mountd(1M)	Answers file system mount requests by reading `/etc/exports` to determine which file systems are available to nodes and users; invoked by `inetd`. The `showmount` command calls `rpc.mountd` to list the clients with local file systems mounted. If operating in an HP-UX cluster environment, `mountd` must be running on any cluster cnode that wishes to export its local file system via NFS. The `mountd` servers are ignored on any cluster client that does not have locally mounted file systems.
rstatd(1M)	Returns statistics obtained from the kernel; invoked by `inetd`. The `rup` program uses `rpc.rstatd`.
rusersd(1M)	Lists the users on the local host; invoked by `inetd`. The `rpc.rusersd` server provides the `rusers` program information about the local users. The `rusers` program then sums and displays the information.
rwalld(1M)	Handles all `rwall` requests; invoked by `inetd`. The RPC program `rwall` sends a message to `rpc.rwalld` on a given host. Each `rpc.rwalld` accepts this message and writes it to all users on the host it is serving using `wall`.
sprayd(1M)	Records the packets sent by `spray`; invoked by `inetd`.

NFS Configuration

Configuring your system is the process of setting up your software so it operates correctly and according to your specifications. The following sections describe the steps you must perform to configure NFS Services on nodes that reside on your network. You can perform some NFS Services configurations in SAM (System Administration Manager), a tool that automates the configuration process. Go to the following sections for detailed configuration instructions (notice that both the SAM and manual configuration methods are included here):

- Compare the files in the `/etc/newconfig` directory to their corresponding existing files.
- Set UIDs and GIDs.
- Create an NFS server and an NFS client using SAM.
- Create an NFS server manually (without SAM).
- Create an NFS client manually (without SAM).
- If applicable, configure the Network Information Service (NIS). (Refer to the "NIS Configuration and Maintenance" chapter.)
- If applicable, configure the Virtual Home Environment (VHE) service. (Refer to "VHE Configuration and Maintenance" chapter.)
- Execute `/etc/netnfsrc` (or reboot) when you are finished with all of the configuration, including setting up NIS and VHE.

Compare /etc/newconfig Files to Existing Files

When you installed the NFS Services software, several new files were copied into the `/etc/newconfig` directory. Perform the following steps to prepare the NFS Service for configuration:

1. Compare each `/etc/newconfig` file listed below with its counterpart shown in the following list.

File in /etc/newconfig directory	Counterpart in /etc directory
netgroup	netgroup
netnfsrc	netnfsrc
rpc	rpc
netnfsrc2	netnfsrc2

2. If the files are the same, then skip to the next section, "Set UIDs and GIDs."

3. If you have previously customized the files that exist in the /etc directory or if the files are from an older version of the software, they will differ from those in /etc/newconfig. If there are differences, copy the current files in /etc to a safe location and do *one* of the following:
   - Change the versions in /etc to reflect the differences in the files in /etc/newconfig.
   - Copy the files in /etc/newconfig to /etc. Then customize the files in /etc if necessary.

Note For the 8.0 release, /etc/netnfsrc has changed significantly. For this reason, /etc/update places the new file in the /etc directory and saves the old file in /etc/netnfsrc.OLD.

Set UIDs and GIDs

The UID field from an /etc/passwd entry and the GID field from an /etc/group entry authenticate NFS users. The client passes this UID and GID to a server for use when checking file ownership and permission.

To ensure only the users in the correct group receive the privileges set by the file's owner, edit /etc/passwd and /etc/group so that each user has one unique UID and one unique GID that is the same on all servers and clients.

If you are using the Network Information Service (NIS), you can configure NIS so you can centrally administer /etc/passwd and /etc/group. (Local UIDs and GIDs are not required if you are using NIS.)

If you are not using NIS, you can use *one* of the following two methods to either create new /etc/passwd and /etc/group files or modify the existing ones:

- Create one /etc/passwd and one /etc/group file to ensure UIDs and GIDs are consistent for each NFS user across the network. Copy these files to all NFS network nodes.

  When updating UIDs or GIDs, you will need to recopy the files to each node. You can automate this process by using shell scripts and the ARPA Services.

  A disadvantage of this method is that it gives exactly the same access to all users across the network. A user with a valid password for a superuser account would have superuser privileges on all nodes configured in this fashion.

- Edit /etc/passwd and /etc/group on each node to ensure UIDs and GIDs are consistent for each user across the network.

  If you modify UIDs or GIDs affecting more than one node, you will have to modify each node affected by the change. For example, if adding a new user you will need to update the /etc/passwd and /etc/group files residing on each system to which the new user will have access.

  Though more time consuming and error prone, this method allows each system to have a different set of users.

Create an NFS Server and an NFS Client Using SAM

SAM (System Administration Manager) provides an automatic method for configuring your local system to be an NFS server or NFS client. You must be superuser to use SAM. The steps covered in this section include:

- Tips for using SAM.
- Getting to the `NFS (Network File System) Configuration` menu.
- Add or modify the connectivity information about a remote system (edit `/etc/hosts` and possibly add `/etc/route` entries to `/etc/netlinkrc`).
- Specify the default gateway.
- Allow this system to access remote file systems via NFS (become an NFS client).
- Add (mount) an NFS file system (edit `/etc/checklist`).
- Allow remote systems to access local file systems via NFS (become an NFS server).
- Modify which systems can access local file systems (edit `/etc/exports`).
- Modify RPC (Remote Procedure Call) services' security (edit `/usr/adm/inetd.sec`).
- Reboot using SAM.

Tips for using SAM

Remember the following tips when you use SAM. You can also get more information from the main menu item, `How to Use SAM`.

- Use your keyboard's cursor control and editing keys to navigate and edit forms.
- You can select a menu item using *either one* of the following methods:
  - Move the cursor to the menu item using [Tab] , the space bar, or the arrow keys and press [Return] or Select Item .
  - Type enough of the menu item's first word to uniquely identify it. In some cases, this is simply the first letter of the menu item. This method does not work for menu items that start with the same word.

- Access the on-line help screens whenever you need more information, such as how or where to obtain a required configuration value! Note that the RESULT sections of the on-line help screens explain what SAM will do "behind the scenes," such as what files SAM will create or modify, or what commands SAM will execute automatically.

Move to the NFS Configuration Menu

All NFS configurations available in SAM are done in the NFS Configuration menu. This section explains how to move to the NFS configuration menu where you can select the task you wish to perform.

1. At the HP-UX prompt, type:

   ```
   sam
   ```

   Wait for SAM's main menu to appear.

2. Select `Networks/Communications`.

   A `WORKING` prompt will appear momentarily while the system is loading the `Networks/Communications` menu.

3. Select `LAN Hardware and Software (Cards and Services)`.

   NFS is not supported over the X.25 link product, so do not choose the X.25 selection.

4. Select `NFS (Network File System) Configuration`.

   A `WORKING` prompt will appear momentarily while the system is loading the menu.

Note You can also get to the NFS configuration menu by selecting `File Systems` from the main menu.

This is the menu where you configure NFS Services. Notice the highlighted *Details* information at the bottom of the screen. This information briefly describes what SAM does if you select the menu item. This information changes as you move up and down the menu.

Note You cannot configure `/etc/netgroup` in SAM. If you wish to edit this file, go to the section, "Create an NFS Server Manually."

Add or Modify Connectivity Information about a Remote System (Edit /etc/hosts and Possibly /etc/netlinkrc)

This task selection allows you to edit `/etc/hosts` and possibly add `/etc/route` entries to `/etc/netlinkrc`, depending on your entries.

Note the following information before you begin:

- If your system is configured to use the NFS Network Information Service (NIS) or the ARPA Service's BIND Name Service for hostname-to-address mapping, you cannot use SAM to add NFS Services connectivity information about a remote system. The `Add/Modify Connectivity Info About a Remote System` form edits only the `/etc/hosts` file; it does not edit an NFS Network Information Service or BIND Name Service database.
- If you must go through a gateway to reach the remote system you are adding connectivity information about, SAM will prompt you for the gateway's hostname and IP address. With this information, SAM will automatically configure the necessary routing by executing an `/etc/route add host` command and adding it to `/etc/netlinkrc`.
- If there is just one gateway you use to reach all systems on other parts of the network, use the `Specify the Default Gateway` form to avoid having to enter the same gateway information every time SAM prompts you for it.

Information you need to complete this task includes:

- Official host name of the remote system.
- IP Address of the remote system.
- Alias names (optional).

To perform this task:

1. Select `Add/Modify Connectivity Info About a Remote System`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. Press **Perform Task**.
4. To verify your entry, press **Exit Task**, then select `View/Remove Connectivity Info About a Remote System`. Follow the instructions to view the entry you just made.
5. If you have another task to perform, press **Exit Task** to go back to the NFS configuration menu. If you have no further tasks to perform, press **Main Menu** and **Exit SAM** to exit SAM.

Specify the Default Gateway

This task selection allows you to set up the default gateway. The only information you need to complete this task is the host name of the default gateway. However, the connectivity information about the gateway must already have been added.

To perform this task:

1. Select `Specify the Default Gateway`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. Press **Perform Task**.
4. If you have another task to perform, press **Exit Task** to go back to the NFS configuration menu. If you have no further tasks to perform, press **Main Menu** and **Exit SAM** to exit SAM.

Allow This System to Access Remote File Systems via NFS (Become an NFS Client)

This task selection allows you to access remote file systems from your local system. When you perform this task you are editing the `/etc/netnfsrc` file. To perform this task:

1. Select `Allow This System to Access Remote File Systems via NFS.`

   If the NFS configuration menu reads `Prevent This System from Accessing Remote File Systems via NFS`, then your system is already set up to be a client. Skip to the next task you wish to perform.

2. Answer "y" (yes) to the question in the pop-up window to allow your system to be an NFS client.

Go to the next section to add the remote file systems you wish to access.

Add (Mount) an NFS File System

This task allows you to add the NFS file systems you wish to access. When you perform this task, you are editing the `/etc/checklist` file. The information you need to complete this task includes:

- Remote system name – system name where the file system you wish to access resides.
- Remote mount directory – directory name of the file system you wish to access.
- Local mount directory – local directory where you want the remote file system be mounted.

To perform this task:

1. Select `Add (Mount) an NFS File System.`
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. Press **Perform Task**.
4. If you have another task to perform, press **Exit Task** to go back to the NFS configuration menu. If you have no further tasks to perform, press **Main Menu** and **Exit SAM** to exit SAM.

Allow Remote Systems to Access Local File Systems via NFS (Become an NFS Server)

This task allows remote systems to have NFS access to your local file systems. When you perform this task you are editing the `/etc/netnfsrc` file. To perform this task:

1. Select `Allow Remote Systems to Access Local File Systems via NFS`.

   If the NFS configuration menu reads `Prevent Remote Systems from Accessing Local File Systems via NFS`, then your system is already set up to be a server. Skip to the next task you wish to perform.

2. Answer "y" (yes) to the question in the pop-up window to allow your system to be an NFS server.

Note When you use SAM to set up your system as an NFS server, it also becomes a PC-NFS server.

Go to the next section to determine which remote systems can have NFS access to your local file systems.

View or Modify Which Systems Can Access Local File Systems

This task allows you to modify a client's NFS access of your local file systems. When you perform this task, you are editing the `/etc/exports` file:

1. Select `View/Modify Which Systems Can Access Local File Systems`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. Press **Perform Task**.
4. If you have another task to perform, press **Exit Task** to go back to the NFS configuration menu. If you have no further tasks to perform, press **Main Menu** and **Exit SAM** to exit SAM.

View or Modify Remote Procedure Call (RPC) Services' Security

This task lets you allow or deny access to specific RPC services (servers). When you perform this task, you are editing the `/usr/adm/inetd.sec` file. The information you need to complete this task are the remote system names to which you are allowing or denying access. To perform this task:

1. Select `View/Modify RPC (Remote Procedure Call) Services' Security`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. If you need to modify more RPC services' security, follow the instructions on the form.
4. When you are finished, press **Perform Task**.
5. If you have another task to perform, press **Exit Task** to go back to the NFS configuration menu. If you have no further tasks to perform, press **Main Menu** and **Exit SAM** to exit SAM.

Rebooting in SAM

If you perform a task that requires a system reboot, a message appears to let you know that a kernel regeneration and reboot are necessary. When this happens:

1. Finish all your NFS configuration tasks.
2. Press **Main Menu** and **Exit SAM** to exit SAM.
3. When the kernel regeneration and reboot message appears, select kernel regeneration.
4. When the kernel regeneration is complete and the reboot message appears, reboot your system.

Caution

Before rebooting, be sure no one is logged onto your system . If you *do not* reboot the system when you exit SAM, NFS Services will not run with the configurations you just made.

If you are configuring NFS in an HP-UX cluster, you must configure NFS into the kernels of all nodes in the cluster. To do this using SAM, run SAM on each node and select `Allow This System to Access Remote File Systems via NFS`.

Create an NFS Server Manually (Without SAM)

You must be superuser to create an NFS server. To create an NFS server, complete the following steps. These steps are described in detail in the sections that follow.

1. Edit `/etc/netnfsrc`.
2. Edit `/etc/inetd.conf`.
3. Edit `/usr/adm/inetd.sec` (if necessary).
4. Edit `/etc/hosts`.
5. Edit `/etc/netgroup` (optional).
6. Create and Edit `/etc/exports`.
7. Reboot the system (if necessary).

An NFS server can also be configured as any combination of an NFS client, NIS client, or NIS server. (An NIS server *must* also be configured as an NIS client.)

Note If you are configuring NFS in an HP-UX cluster environment, you must configure NFS into the kernel on all cnodes in the cluster. See "Configure a New Kernel" in Chapter 3.

---

1. Edit /etc/netnfsrc

The `/etc/netnfsrc` file activates the NFS daemons and servers.

- To define the node as an NFS server, set the `NFS_SERVER` variable to any digit other than zero.
- If the node is also a client, you may want to set the `NFS_CLIENT` variable to any digit other than zero now. (Refer to the "Create an NFS Client Manually" section to complete client configuration procedures.)
- If the node is also a server for PC-NFS requests, set the `PCNFS_SERVER` variable to any digit other than zero.

Client Only	NFS_CLIENT=1 NFS_SERVER=0
Server Only	NFS_CLIENT=0 NFS_SERVER=1
Both Client and Server	NFS_CLIENT=1 NFS_SERVER=1
Neither Client nor Server	NFS_CLIENT=0 NFS_SERVER=0
PC-NFS Server	PCNFS_SERVER=1

You can refer directly to the comments (lines beginning with pound signs) for editing instructions and for descriptions of each activity executed by `/etc/netnfsrc`.

Note If you edit this file other than specified in this document, HP recommends you incorporate personal comments for future system administration.

```
#!/bin/sh
#       netnfsrc        NFS startup file
##
#       Depending on the configuration parameters you set within,
#       this script sets up some or all of the following:
#*      NIS specific:
#           domainname          the NIS domain name
#
#       and starts up some or all of the following programs:
#           portmap             RPC (program_#,version) -> port_# mapper
#           nfsd                NFS daemons
#           biod                async BIO daemons
#           pcnfsd              PC-NFS daemon
#*      NIS specific:
#           ypbind              NIS client process (all NIS nodes)
#           ypserv              NIS server process (NIS server only)
#           yppasswdd           NIS password daemon (NIS master server only)
##
#       NFS_CLIENT              1 if this node is an NFS client, 0 if not
#       NFS_SERVER              1 if this node is an NFS server, 0 if not
#       Note:   it is possible for one host to be a client, a server, both
#               or neither!  This system is an NFS client if you will be
#               NFS mounting remote file systems; this system is a server
#               if you will be exporting file systems to remote hosts.
#       See Also: nfsd(1M), mount(1M)
##
NFS_CLIENT=0
NFS_SERVER=0
.
.
.
.
PCNFS_SERVER=0
```

2. Edit /etc/inetd.conf

To activate the RPC services, remove all # comment marks (pound signs) from `/etc/inetd.conf` lines beginning with `#rpc`. If you want one of these services activated but the line was removed, you may need to obtain a new version of `/etc/inetd.conf` from `/etc/newconfig`.

Note After editing `/etc/inetd.conf`, you must reconfigure `inetd` by entering:

```
/etc/inetd -c
```

RPC Services Security. The `inetd` security facility works only when the `inetd` executes a server. For the RPC services that do not exit after each service request, `inetd` provides a security check only for the first request. Successive requests bypass the `inetd` and are subject only to the security checking performed by the individual RPC services. However, you can make the `inetd` perform a security check for every RPC request by doing *both* of the following steps:

- Add the `-e` option to the `/etc/inetd.conf` entry for the RPC service.
- Specify the RPC service in the first field of `/usr/adm/inetd.sec`. (Refer to the next section, "Edit `/usr/adm/inetd.sec`.")

Note Adding the `-e` option makes the RPC server respond slower since it has to restart for each request.

For information on C2 Security, refer to the *HP-UX System Security Manual* and the *HP-UX Beginner's Guide*.

RPC Entries. Refer to the following list for a brief description of each RPC service line present in `/etc/inetd.conf`.

```
rpc dgram udp wait root /usr/etc/rpc.mountd 100005 1 rpc.mountd -e
```

The rpc.mountd program is the server for the mount command and reads /etc/exports to see what the available file systems are and to whom they are exported. It also keeps a list of all mounted file systems. The program supports version 1.

The -e option forces inetd to perform a security check for rpc.mountd on every request.

```
rpc stream tcp nowait root /usr/etc/rpc.rexd 100017 1 rpc.rexd
```

The rpc.rexd program is the server for the on program. The program supports version 1.

```
rpc dgram udp wait root /usr/etc/rpc.rstatd 100001 1-3 rpc.rstatd
```

The rpc.rstatd program is the server for the rup command and provides kernel statistics. The program supports versions 1 through 3.

```
rpc dgram udp wait root /usr/etc/rpc.rusersd 100002 1-2 rpc.rusersd
```

The rpc.rusersd program is the server for the rusers command and provides information about active users on remote nodes and the amount of time they have been idle. The program supports versions 1 and 2.

```
rpc dgram udp wait root /usr/etc/rpc.rwalld 100008 1 rpc.rwalld
```

The rpc.rwalld program writes a message sent by rwall to all users logged on to the system. The program supports version 1.

```
rpc dgram udp wait root /usr/etc/rpc.sprayd 100012 1 rpc.sprayd
```

The rpc.sprayd program is the server for the spray command and accepts RPC requests, reads UDP packets, and then tells how fast it read them; you can use the results to gauge performance. The program supports version 1.

```
rpc dgram udp wait root /usr/etc/rpc.rquotad 100011 1 rpc.rquotad
```

The rpc.rquotad program is the server for the quota command. The daemon returns data regarding disk quotas for NFS mounted file systems.

3. Edit /usr/adm/inetd.sec (if necessary)

NFS operates under the assumption you have a "friendly" network; meaning, you can trust all users attached to your network. Since this assumption may not apply to everyone, refer to the following sections to improve your file security.

The `/usr/adm/inetd.sec` configuration file is provided in the ARPA Services product. It is not solely for NFS access.

This file allows you to determine:

- How many remote services can run simultaneously on the local host.
- Which hosts are allowed to remotely use the local host.

Note If `inetd` is running, it rereads `/usr/adm/inetd.sec` after you make changes to it. Your changes apply only to services started after the file is reread, but not to any currently running services.

Set Maximum Number of Remote Connections. On the first line in `/usr/adm/inetd.sec`, enter the maximum number of simultaneous remote services to be started by `inetd` as shown in the following example:

```
MAXNUM number
```

If you do not specify a `MAXNUM` value, the default is 1000.

Specify Accesses to Services. Each entry in `/usr/adm/inetd.sec` has the following format (enter either `allow` or `deny`):

```
service_name allow/deny host_specifier(s)
```

/usr/adm/inetd.sec Entry Fields	Description
service_name	Name of a valid service (including RPC services) with an entry in /etc/inetd.conf. - For RPC services, *service_name* is the name of the service that matches its program number in /etc/rpc. This entry *must* have a corresponding entry in /etc/inetd.conf which contains the -e option. - Specify only one service per entry. - If an entry in /usr/adm/inetd.sec specifies the service name and nothing else, inetd allows all hosts to attempt access.
allow/deny	The *allow* entry instructs inetd to approve the host or network for access to the specified service. The deny entry instructs inetd to disapprove the host or network for access to the specified service.
host_specifier(s)	Name of a host or a network listed in /etc/hosts or /etc/networks, or an internet address in the standard internet notation. - You can specify more than one host or network by separating each *host_specifier* with a blank or tab. - You can use the * (wild card character) or - (range character) in any field of a network or host address. - You cannot use aliases.

RPC Services Security. You can make `inetd` perform its `inetd.sec` security check for every RPC request by following these two steps.

1. Add the `-e` option to the RPC service line in `/etc/inetd.conf`. (Refer to the "2. Edit /etc/inetd.conf" section or `inetd.conf`.)

   EXAMPLE: `rpc dgram udp wait root /usr/etc/rpc.mountd 100005 1 rpc.mountd -e`

2. Specify the RPC service in the first field in `/usr/adm/inetd.sec`.

/usr/adm/inetd.sec Example RPC Entry	Effect on System Security
`mountd allow hostA`	Allows only `hostA` to access `rpc.mountd`
`walld deny 111.56.78.9 10.*`	Denies access to `rpc.rwalld` from the following hosts: - 111.56.78.9 (internet address) - all hosts that are part of network 10.*

4. Edit /etc/hosts

Caution If NIS is running, do not edit `/etc/hosts` on any node except the NIS master server; otherwise, local changes will not be propagated.

If you have ARPA Services and have configured the BIND name server, do not edit the `/etc/hosts` file. See the "Maintaining Network and Domain Data Files" section of the BIND Name Server chapter in *Installing and Administering ARPA Services*.

As node manager, you must configure this file for your host. You can add entries to this file either automatically with the System Administration Manager (SAM) or manually by editing the file. This section describes how to configure `/etc/hosts` manually.

Adding IP Addresses. If your host has more than one IP address (for multiple network interfaces), you must add entries for every IP address. These entries must have the same

official host name but different aliases or have different official hostnames. This is so that different IP addresses (network interfaces) are distinguished (and can be referenced) by different aliases s or hostnames.

Note You can copy the official host data base maintained at the Network Information Control Center (NIC) for ARPA Internet networks. (Refer to the "Military Standards and Request for Comment Documents" section of Chapter 1 for information on how to contact the NIC.) Be sure to check the format of files received from the NIC.

If your host accesses a multi-homed host (one with more than one link interface), make sure the internet address for that host is correct in the `/etc/hosts` file with respect to your host. For example, in the networking scheme Figure 4-1, host `paul` and host `barb` access multi-homed host `mickie` via internet address 192.6.21.2. Hosts `dean` and `dennis`, on the other hand, access host `mickie` via internet address 192.6.36.3.

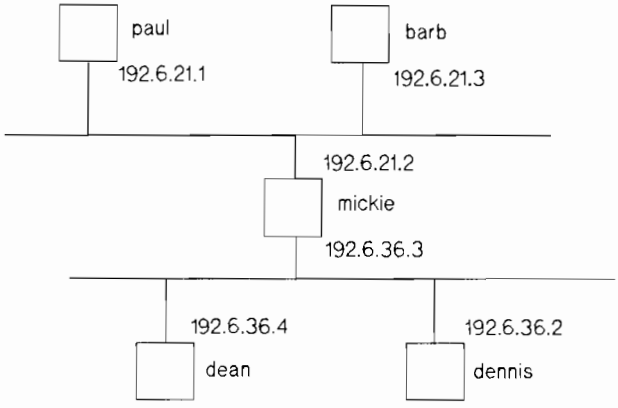

Figure 4-1. Multi-Homed Host Network Scheme

Syntax for /etc/hosts. Each host (including the local host) has a one line entry in the `/etc/hosts` file. Each entry in the `/etc/hosts` file takes the following form:

internet_address official_host_name [alias(es)]

internet_address Network address that uniquely identifies the node. *Internet_address* must be in dot notation. Refer to the "Assigning an Internet Address" section in *Installing and Administering LAN* for more information on internet addresses.

official_host_name Name of the node. Host names can contain any printable character except white spaces, newline, or the comment character (#). By convention, *official_host_name* should be the same as the system host name assigned with the HP-UX `hostname` command.

alias(es) Common name or names for the node. An *alias* is a substitute for *official_host_name*. *Alias* names are optional and are not supported by all the commands that use `/etc/hosts`.

Format for /etc/hosts

- Lines cannot start with a white space (tabs or blanks).
- The fields can have any number of blanks or tab characters separating them.
- Comments are allowed and are designated by a "#" in front of the comment text.
- Trailing blanks and tab characters are allowed.
- Blank lines are allowed.

Example of /etc/hosts Entry. The `/etc/hosts` entry for a node with:

- The address 192.45.36.5.
- The official host name *hpdxsg*.
- The alias name *bullfrog*.

Looks like:

```
192.45.36.5  hpdxsg  bullfrog
```

Permissions. The `/etc/hosts` file should be owned by user *root*, group *other* and have 0444 (r – r – r –) permission.

Refer to the /etc/hosts file for examples of the actual format and contents. For more information on /etc/hosts, refer to the *hosts*(4) entry in the *HP-UX Reference*.

Verification. To view the list of remote systems you may communicate with, type the following command at the HP-UX prompt:

```
more /etc/hosts
```

To verify that /etc/hosts is being used to do host name to address mapping, use nslookup as described in the previous section, "3. Configure Host Name to Address Mapping."

To view the destinations reached through gateways and the gateways used to reach those destinations, type the following command at the HP-UX prompt:

```
netstat -r
```

The listing from this command may appear slowly, as it attempts to find the names associated with the network addresses used to perform routing.

Copying a Remote /etc/hosts File to Your Local Host. When you first configure your host's /etc/hosts file, it is very small. If you want to get a copy of a larger, more complete /etc/hosts file from another host, you can do it *one* of two ways:

- If you have ARPA Services on your host, go to the *Installing and Administering ARPA Services* manual and use the method described in the "Editing /etc/hosts" section of Chapter 2.
- If you have NFS Services ONLY, the method is more complicated and is described in the following example.

In the following example, your local host is named myhost and the remote host that has the complete /etc/hosts file is named otherhost. Perform the following steps:

1. Using either SAM or the manual method, add connectivity information about otherhost to the /etc/hosts file on myhost.
2. Using either SAM or the manual method, make otherhost an NFS server so that it allows myhost access to the otherhost root (/) file system.
3. Using either SAM or the manual method, make myhost an NFS client.

4. On myhost, mount the otherhost file system "/", copy /etc/hosts from otherhost, then unmount the otherhost "/" file system. See the following example:

```
mkdir /tmp/exmpl
mount otherhost:/ /tmp/exmpl
cp /tmp/exmpl/etc/hosts /etc/hosts
umount /tmp/exmpl
rmdir /tmp/exmpl
```

If you overwrite your local /etc/hosts file with a copy from another host, you may need to bring it up to date by adding unofficial aliases or unknown hosts, including your own host.

5. Edit /etc/netgroup

Caution If NIS is running, do not edit /etc/netgroup on any node except the NIS master server; otherwise, local changes will not be propagated.

The /etc/netgroup file enables you to define a specific network-wide group of nodes as a **netgroup**. You can then limit file system access by exporting file systems (via /etc/exports) to the netgroups defined.

The system uses /etc/netgroup to verify host names whenever clients perform remote mounts. (Refer to *netgroup*(4) in the *HP-UX Reference*.)

For *ARPA Services*, the system uses /etc/netgroup to verify users when clients perform remote logins or remote shells. (Refer to *hosts.equiv*(4) in the *HP-UX Reference*.)

Add a line with the following format for each netgroup you wish to define.

The entry may contain any number of netgroup names:

```
netgroup_name1  netgroup_name2  netgroup_name3  ...
```

But then you must define these netgroups within `/etc/netgroup`:

```
netgroup_name1  member1  member2  ...
```

You can use the following conventions when editing the `/etc/exports` file:

- The member*n* is equal to the triple (*host_name*, *user_name*, *NIS_domain_name*).
- You can assign more than one triple to a netgroup by enclosing each separate set within parentheses (*host_name*, *user_name*, *NIS_domain_name*).
- Leave any of these three fields empty to signify a wild card (i.e., blank fields match anything). For example, `(,,research)` matches all hosts and users in the `research` NIS domain.
- A - (dash) in any of these three fields means *match nothing*. For example, `(-,mike,graphs)` does not match any hosts, but it does match the user *mike* in the `graphs` NIS domain.
- Each *host_name* must have an entry in `/etc/hosts`.
- The *NIS_domain_name* is the name of the NIS domain to which you currently belong. To determine your current NIS domain name, execute the `domainname` command.

  The commands using `/etc/netgroup` assume you are not looking for any NIS domain other than the one assigned on your node.

EXAMPLES:

/etc/netgroup Example Entry	The Netgroup Includes
`netgroup1 (,,)`	Everyone on the network.
`netgroup2 (,darren,graphic)`	The user `darren` on any host in the `graphic` NIS domain.
`netgroup3 (node_7,,graphic)`	Any user on the `node_7` host in the `graphic` NIS domain.
`netgroup4 (node_2,john,)`	The user `john` on the `node_2` host in any NIS domain.
`netgroup5 (,andy,graphic) (node_1,mike,)`	The user `andy` on any host in the `graphic` NIS domain and the user `mike` on the `node_1` host in any NIS domain.
`netgroup6 (-,annette,graphic)`	The user `annette` in the `graphic` NIS domain, no host included.

6. Create and Edit /etc/exports

You control the available file systems by your entries in the server's `/etc/exports` file. Each time a server receives a mount request, `rpc.mountd` accesses `/etc/exports` to see if the file system is exported and which systems can access it.

- The server must have the file system mounted locally before it can be exported.
- You must export the entire file system; you cannot export specific directories (though clients can mount specific directories).
- The path name in `/etc/exports` must be the same path name as the directory on which the local file system is mounted.
- If the `-async` option is set for a file system, asynchronous writes on the NFS server occur. (See *exports*(4) in the *HP-UX Reference*.)

Caution

The `-async` option increases write performance on the NFS server by allowing asynchronous writes on the NFS server's file system.

However, use caution in deciding whether to use the `-async` option. An unreported data loss may occur if the option is set and the NFS server hardware experiences a power loss, system panic or other failure.

Do not use the `-async` option with file systems that contain:

- Files which are accessed by the `O_SYNCIO` flag (which is set by the `fnctl` or `open` calls).

- Data that cannot be reconstructed (e.g., a file system containing database files).

- Files synchronized with `fsync`.

- Critical applications requiring absolute data integrity.

If you are unsure whether any of the previous conditions apply, do not use the `-async` option.

You control the file system's availability by specifying a netgroup or host name; otherwise, the file system becomes available to everyone on the network running NFS. After accessing `/etc/exports`, the system checks `/etc/netgroup` for the netgroup definition; if it is not present, the system checks `/etc/hosts` for the host name. (For more information, refer to the previous sections, "4. Edit /etc/hosts" and "5. Edit /etc/netgroup.")

Note

If importing a file system containing a user's home directory, the user may not be able to login if the remote file system is not accessible.

If a client has a file system mounted and you edit `/etc/exports` to change availability of that file system, the client's access will not change. To prevent the client from accessing the server's files, on the client you must either unmount the file systems or reboot the client.

/etc/exports Entry Formats	System Response
/complete_filesystem_pathname	Exports the file system to everyone on the network and defaults to *synchronous* writes on the NFS server.
/complete_filesystem_pathname netgroup_1 netgroup_2	Exports the file system only to specified netgroups.
/complete_filesystem_pathname client_1 client_2	Exports the file system only to specified clients.
/complete_filesystem_pathname client_2 netgroup_1	Exports the file system only to the specified client and netgroup.
/complete_filesystem_pathname -async client_1	Exports the file system to the specified client and causes *asynchronous* writes on the NFS server.

EXAMPLE:

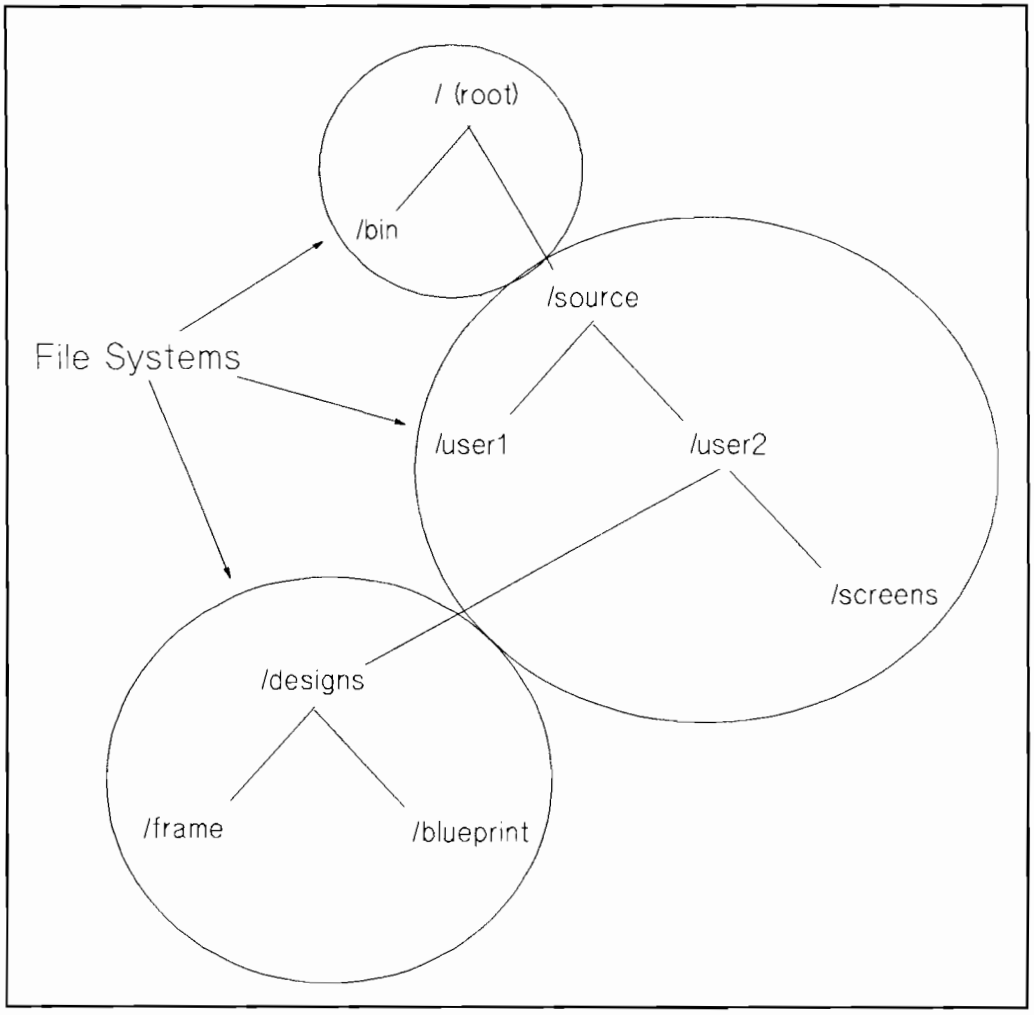

Figure 4-2. /etc/exports Example

/etc/exports Example Entry *	Example System Response
/	Export / (root) to all clients. Clients will not receive the file system `source` since / is the exported file system. The `source` and `designs` file systems will not be seen by this mount. Clients will receive the `/bin` directory since is it part of the / file system.
`/source`	Export `source` to all clients. Clients will not receive the file system `designs` or / since `source` is the exported file system. The `designs` and / file systems will not be seen by this mount.
`/source/user2/designs`	Export `designs` to all clients.
`/source/user2/designs -async system1`	Export `designs` to the client `system1` and allow **asynchronous** writes on the NFS server.
`/source/user2/designs lab`	Export designs to the netgroup `lab`.
`/source/user2/designs system1 lab`	Export `designs` to the client `system1` and the netgroup `lab`.
**Note*: You must define all hosts in `/etc/hosts` and all netgroups in `/etc/netgroup` or, if you are using NIS, ensure that all hosts and netgroups are defined on the NIS master server.	

7. Reboot the System (if necessary)

After you finish the configuration procedure, execute `/etc/netnfsrc` or reboot the system to activate the daemons and servers.

The rebooting process does not unmount any of the server's file systems that were remotely mounted by other network nodes. However, these nodes will not be able to access any of the server's files until the server is operating again.

Create an NFS Client Manually (Without SAM)

You must be superuser to create an NFS client.

To create an NFS client, complete the following steps:

1. Edit `/etc/netnfsrc`.
2. Mount file systems
3. Reboot the system (if necessary).

An NFS client can also be configured as any combination of an NFS server, NIS client, or NIS server. (An NIS server *must* also be configured as an NIS client.)

1. Edit /etc/netnfsrc

The `/etc/netnfsrc` file activates the NFS daemons and servers. Do the following:

- To define the node as an NFS client, set the `NFS_CLIENT` variable to any digit other than zero.
- If the node is also a server, you may want to set the `NFS_SERVER` variable to any digit other than zero now. (Refer to the "Create an NFS Server" section to complete server configuration procedures.)
- If the node is also a server for PC-NFS requests, set the `PCNFS_SERVER` variable to any digit other than zero.

Client Only	`NFS_CLIENT=1` `NFS_SERVER=0`
Server Only	`NFS_CLIENT=0` `NFS_SERVER=1`
Both Client and Server	`NFS_CLIENT=1` `NFS_SERVER=1`
Neither Client nor Server	`NFS_CLIENT=0` `NFS_SERVER=0`
PC-NFS Server	`PCNFS_SERVER=1`

You can refer directly to the comments (lines beginning with # (pound) signs) for editing instructions and for descriptions of each activity executed by `/etc/netnfsrc`.

Note If you edit this file other than specified in this document, HP recommends you incorporate personal comments for future system administration.

```
#!/bin/sh
#       netnfsrc        NFS startup file
##
#       Depending on the configuration parameters you set within,
#       this script sets up some or all of the following:
#*      NIS specific:
#           domainname          the NIS domain name
#
#       and starts up some or all of the following programs:
#           portmap             RPC (program_#,version) -> port_# mapper
#           nfsd                NFS daemons
#           biod                async BIO daemons
#           pcnfsd              PC-NFS daemon
#*      NIS specific:
#           ypbind              NIS client process (all NIS nodes)
#           ypserv              NIS server process (NIS server only)
#           yppasswdd           NIS password daemon (NIS master server only)
##
#       NFS_CLIENT              1 if this node is an NFS client, 0 if not
#       NFS_SERVER              1 if this node is an NFS server, 0 if not
#       Note:   it is possible for one host to be a client, a server, both
#               or neither!  This system is an NFS client if you will be
#               NFS mounting remote file systems; this system is a server
#               if you will be exporting file systems to remote hosts.
#       See Also: nfsd(1M), mount(1M)
##
#       Note:   this has nothing to do with whether or not the system is
#               a rootserver or diskless client workstation.  There is a
#               test for this later.
##
NFS_CLIENT=0
NFS_SERVER=0
.
.
.
.
PCNFS_SERVER=0
```

2. Mount File Systems

Review the servers' `/etc/exports` files on your LAN to determine the file systems to which you want the client to have access. You will need to mount each of these file systems on the clients.

For each file system you should determine *one* of the following mounting methods:

- Mount automatically at boot time via `/etc/checklist`.
- Mount only when manually specified via the `mount` command.

Since an attempt to mount a remote file system requires using another node and the network, the mount may not succeed the first time. You can vary the number of times NFS attempts to mount a file system by using the `retry` option.

After the mount is successful, the manner in which NFS handles requests depends on whether the mount is hard (default) or soft.

NFS Hard Mount

Hard mounted file systems with the default `int` (interrupt) cause NFS to retry a request until it succeeds, you interrupt it, or you reboot the system. If the `noint` option is activated and an NFS server goes down, the system retries the request until the server comes up again or you reboot the system.

If the server does not respond to a hard mount request, NFS writes the following message in the network error log file.

NFS: server `host_name` not responding, still trying

Refer to *Installing and Administering LAN/9000* documentation for more error log information.

Note If a server that you previously performed a hard mount goes down, you may not be able to access mounted file systems on other nodes unless you reboot the problem server or interrupt all its requests.

NFS Soft Mount

Soft mounted file systems abort requests after one attempt. NFS writes an error to the log file if the server does not respond to a request. The message varies depending on what type of request is made.

```
NFS server host_name not responding, giving up
```

```
NFS function_name failed for server server_name: TIMED OUT
```

Note

If a user's home directory is in a remote file system, the user will not be able to login if the remote file system is not accessible (e.g., the server goes down, the network fails).

Mount Guidelines

Refer to the following guidelines whether mounting file systems automatically via the `/etc/checklist` file or manually via the `mount` command. For more specific information, refer to *checklist*(4) and *mount*(1M) in the *HP-UX Reference*.

- You cannot mount a remote file system unless the server has an entry for your node in `/etc/exports` or unless `/etc/exports` makes the file system available to everyone on the network. (Execute `showmount` to list mounted file systems.)
- Though a server can export only file systems, you can mount file systems or directories.
- When you mount a new file system on top of a directory already containing files, the directory's files will no longer be accessible unless you execute `umount` to unmount the mounted file system.

  To avoid masking a directory, HP recommends you mount the file system on top of an empty directory.
- You cannot mount or unmount an open directory (a directory in which someone is currently operating).
- You must specify a **mount point** (name of a local directory on which the file system will be mounted).

- If operating in an HP-UX cluster environment:
  - If a cnode mounts a remote file system, all cnodes in the cluster can access the remote file system.
  - If using NFS to mount a file system attached to a cluster, you must use the host name where the file system is locally mounted, as the node name specified in the `mount` command.
  - If a cnode mounts a remote file system, any cnode in that cluster can unmount the remote file system.
  - All mount points must exist on the file systems mounted on the cluster root server. That is, mount points cannot exist on file systems locally mounted on a cluster auxiliary server.
  - If a cnode that mounted a remote file system goes down, all other cnodes in the cluster can still access that remote file system.
- Before mounting a file system, refer to the following table and determine the options you want the mount to have.
  - You must specify an option if mounting via `/etc/checklist`; you do not have to specify an option if mounting via `mount`.
  - You do not have to list options in a specific order; however, you must separate the options with commas (not spaces).

NFS Mount Options	Description
bg	*Background*: If the first request to a remote node's mountd fails, the mount process continues retrying the request in the background.
defaults	*Defaults*: The mount takes all the default options without you having to individually specify them. The defaults are noted within this table by asterisks (*). You only need to specify defaults when mounting via /etc/checklist; the mount command automatically provides the defaults.
fg*	*Foreground*: If the first request to a remote node's mountd fails, the mountd daemon retries the requests in the foreground.
hard*	*Hard Mount*: NFS retries until the request succeeds or you reboot the system. If you are using the int default option, you can interrupt the file system request.
int*	*Interruptable Mount*: You can press an interrupt key to abort an NFS request. (Though the interrupt key is not defined, common ones include [CTRL]-[C] and [BREAK].)
noauto	*No Automatic Mount*: Prevents the file system from being mounted when the mount -a option is executed. You only need to specify noauto when mounting via /etc/checklist.
nointr	*No Interruptable Mount*: You cannot interrupt processes waiting for NFS requests to complete.
nosuid	*No setuid*: You cannot execute files on the remote file system with either the setuid or setgid bits set.
port = *n*	Port = *n* Default *n* = 2049 (the NFS server port) Specifies the UDP port at which the NFS server is contacted. You should not have to reset this value.
quota	*Disk quotas*: Activate disk quotas on the file system.
* = *Default*	

NFS Mount Options	Description
`retrans` = *n*	Retransmit = *n* Default *n* = 4 When NFS sends a request to a remote system, RPC attempts to transmit the request times. If RPC does not receive a response after *n* attempts, soft mounts return an error and hard mounts retry the request.
`retry` = *n*	Retry = *n* Default *n* = 1 The `mount` command retries mounting the file system *n* times; the default is 1. For example, if a `mount` attempt fails once and the default is 1, `mount` tries once more before quitting.
`ro`	*Read Only*: Access rights are *Read Only*.
`rsize` = *n*	Read requests size = *n* Default *n* = 8192 (8K) Specifies the maximum read request size used in communicating with the server.
`rw`*	*Read/Write*: Access rights are *read* and *write*.
`soft`	*Soft Mount*: NFS aborts the request after RPC attempts to transmit the request `n` times (as specified by the `retrans` option).
`suid`*	*setuid*: You can execute programs on the remote file system that have setuid as one of their permissions.

NFS Mount Options	Description
`timeo` = *n*	Timeout = *n* Default *n* = 7 Specifies the initial timeout (in tenths of seconds) for NFS requests. When an NFS request occurs, RPC sends the request, waits 0.7 seconds for a response, and then retries the request. After the initial timeout, the timeout increases by multiples of two each time no response is received. When a specified number of `retrans` retransmissions have been sent with no reply, soft mounts return an error and hard mounts retry the request. *Note*: If performing NFS mounts through a gateway and you see several `server not responding` messages within a few minutes, change the timeout default value (7) to a value of 10 or greater until you stop seeing the message.
`wsize` = *n*	Write size = *n* Default *n* = 8192 bytes (8K) Specifies the maximum write request size used in communicating with the server.
* = *Default*	

Edit /etc/checklist for Automatic Mounts

If you want the file system mounted automatically, add an entry for it in the `/etc/checklist` file. At boot time, `/etc/netnfsrc2` executes `mount -at nfs` to mount all NFS file systems listed in `/etc/checklist`.

Edit `/etc/checklist` to append the hosts and file systems you wish to import using the following format. All of the default options are activated when you specify `defaults`. You must specify either `defaults` or at least one option.

For NFS Hard Mounts via `/etc/checklist`:

server_name:*/imported_filesystem* */mount_point* `nfs defaults 0 0`

or

server_name:*/imported_filesystem* */mount_point* `nfs` [*options*] `0 0`

For NFS Soft Mounts via `/etc/checklist`:

server_name:*/imported_filesystem* */mount_point* `nfs soft`[`,`,*other options*] `0 0`

The `nfs` stands for NFS mounts. NFS ignores the two zeros (0 0), though they must be present.

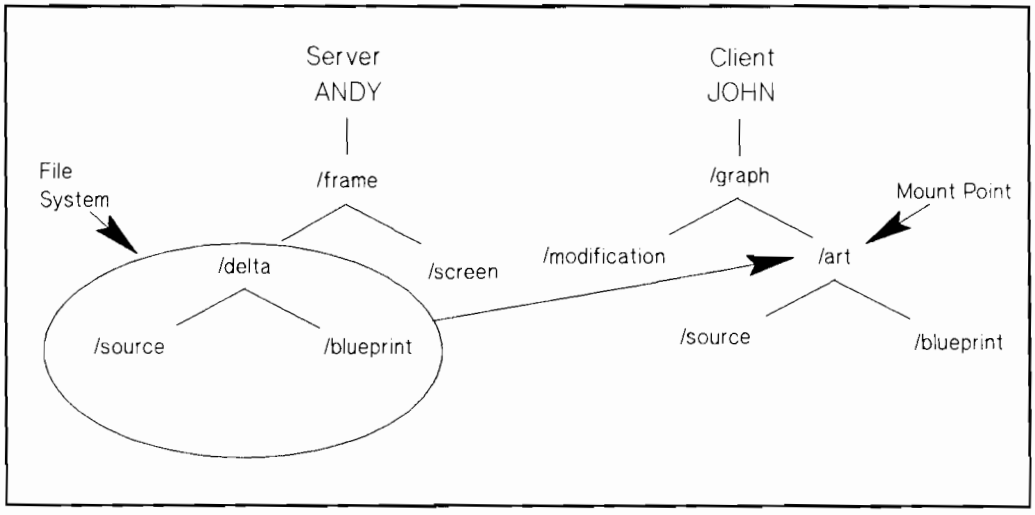

Figure 4-3. /etc/checklist Automatic Mounts

/etc/checklist Example Entry on the Client JOHN	Resulting Mount Options*
ANDY:/frame/delta /graph/art nfs defaults 0 0	Foreground Hard Mount Interruptable Port = 2049 Read and Write Read Size = 8192 Retransmit = 4 Retry = 1 setuid Timeout = 0.7 Write Size = 8192 *Note*: All of these options are by default.
ANDY:/frame/delta /graph/art nfs ro,retry=6,timeo=3 0 0	Read Only Retry = 6 Timeout =0.3
ANDY:/frame/delta /graph/art nfs bg,retrans=8,soft 0 0	Background Retransmit = 8 Soft Mount
ANDY:/frame/delta /graph/art nfs noauto,noint,nosuid 0 0	No Automatic Mount No Interruptable Mount No setuid
ANDY:/frame/delta /graph/art nfs rsize=1024,wsize=1024 0 0	Read Size = 1024 bytes Write Size = 1024 bytes

* The default options are activated when you specify `defaults`. They are also active with other options unless you specify otherwise. The default options are listed only once for this example.

Execute mount for Manual Mounts

Execute `mount` to mount an NFS file system manually. NFS file systems mounted via `mount` are only mounted as long as the client is running or until they are unmounted via `umount`. If the client goes down, you will have to re-mount the file system.

Do not use `mount` if you listed the file system in `/etc/checklist` since it will have already been mounted.

Use the following `mount` format for NFS mounts. All of the default options are activated unless you specify otherwise.

For NFS Hard Mounts via mount:

```
mount [-o options] server_name:/filesystem /mount_point
```

For NFS Soft Mounts via mount:

```
mount -o soft[,other_options] server_name:/filesystem /mount_point
```

EXAMPLES: `mount` Manual NFS Mounts

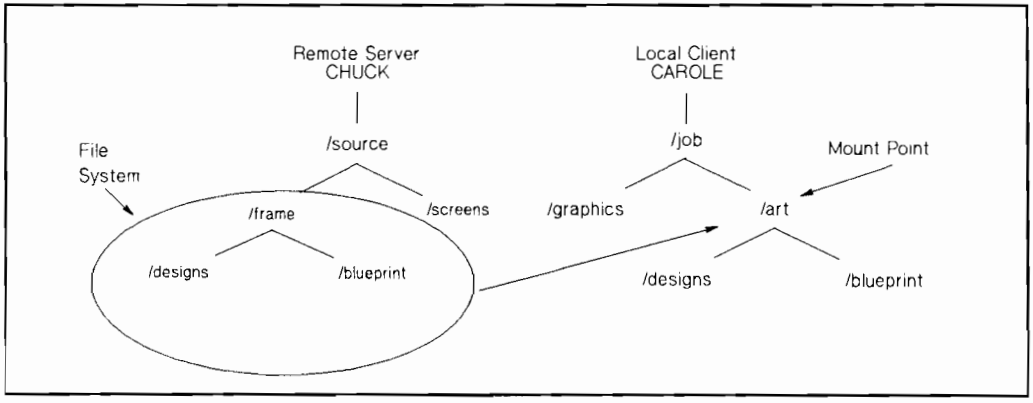

Figure 4-4. Example of a Manual NFS Mount

mount Example Command	Resulting Mount Options*
mount CHUCK:/source/frame /job/art	Foreground Hard Mount Interruptable Port = 2049 Read and Write Read Size = 8192 Retransmit = 4 Retry = 1 setuid Timeout = 0.7 Write Size = 8192 *Note*: All of these options are by default.
mount -o ro,retrans=8,timeo=3 CHUCK:/source/frame /job/art	Read Only Retransmit = 8 Timeout = 0.3
mount -o bg,retry=6,rw,soft CHUCK:/source/frame /job/art	Background Read and Write Retry = 6 Soft Mount
mount -o noauto,noint,nosuid CHUCK:/source/frame /job/art	No Automatic Mount No Interruptable Mount No setuid
mount -o rsize=1024,wsize=1024 CHUCK:/source/frame /job/art	Read Size = 1024 bytes Write Size = 1024 bytes

* All of the default options are activated unless you specify otherwise. The default options are listed only once for this example.

3. Reboot the System (if necessary)

After you finish the configuration procedures, execute `/etc/netnfsrc` or reboot the servers and clients to activate the daemons and servers.

The rebooting process unmounts all local file systems and directories that were manually mounted by the client (i.e., were not automatically mounted by `/etc/checklist`).

Configure NIS (optional)

If you plan to use the optional Network Information Service (NIS), refer to the "NIS Configuration and Maintenance" chapter for detailed configuration procedures.

Configure VHE (optional)

If you plan to use the optional Virtual Home Environment (VHE) service, refer to the "VHE Configuration and Maintenance" chapter for detailed configuration procedures.

Execute /etc/netnfsrc

To complete the configuration procedure, execute `/etc/netnfsrc` (or reboot) your system.

Note You have completed configuring the base NFS service. Refer to the remaining part of the chapter for maintenance information.

NFS Maintenance

To keep NFS running correctly and efficiently, refer to the following sections to ensure it stays configured to meet your changing needs:

- Maintain NFS Services using SAM.
- Prevent systems from accessing local file systems via NFS (without using SAM).
- Update software.
- Clock skew.
- Maintain the NFS server.

Maintain NFS Services Using SAM

SAM (System Administration Manager) provides an automatic method for maintaining your local system's client and server services. You must be superuser to use SAM. The sections covered include:

- Tips for using SAM.
- Move to the `NFS (Network File System) Configuration` menu.
- View or remove connectivity information about a remote system.
- Prevent this system from accessing remote file systems via NFS (stop being an NFS client).
- Modify NFS mount options.
- Remove (unmount) an NFS file system.
- Prevent remote systems from accessing local file systems via NFS (stop being an NFS server).
- View or modify which systems can access local file systems.

Tips for using SAM

Remember the following tips when you use SAM. You can also get more information from the main menu item, `How to Use SAM`.

- Use your keyboard's cursor control and editing keys to navigate and edit forms.
- You can select a menu item using *either one* of the following methods:
  - Move the cursor to the menu item using [Tab] , the space bar, or the arrow keys and press [Return] or Select Item.
  - Type enough of the menu item's first word to uniquely identify it. In some cases, this is simply the first letter of the menu item. This method does not work for menu items that start with the same word.
- Access the on-line help screens whenever you need more information, such as how or where to obtain a required configuration value! Note that the RESULT sections of the on-line help screens explain what SAM will do "behind the scenes," such as what files SAM will create or modify, or what commands SAM will execute automatically.

Move to the NFS Configuration Menu

All NFS configurations available in SAM are done in the NFS Configuration menu. This section explains how to move to the NFS configuration menu where you can select the task you wish to perform.

1. At the HP-UX prompt, type:

   ```
   sam
   ```

   Wait for SAM's main menu to appear.

2. Select Networks/Communications.

   A WORKING prompt will appear momentarily while the system is loading the Networks/Communications menu.

3. Select LAN Hardware and Software (Cards and Services).

   NFS is not supported over the X.25 link product, so do not choose the X.25 selection.

4. Select NFS (Network File System) Configuration.

   A WORKING prompt will appear momentarily while the system is loading the menu.

Note You can also get to the NFS configuration menu by selecting `File Systems` from the main menu.

This is the menu where you configure NFS Services. You cannot configure `/etc/netgroup` in SAM. If you wish to edit this file, go to the section, "Create an NFS Server Manually."

Notice the highlighted *Details* information at the bottom of the screen. This information briefly describes what SAM does if you select the menu item. This information changes as you move up and down the menu.

View or Remove Connectivity Information about a Remote System

This task selection allows you to delete entries from `/etc/hosts` and possibly delete `/etc/route` entries from `/etc/netlinkrc`, depending on the information you are deleting.

Note the following information before you begin:

- If your system is configured to use the NFS Network Information Service or the ARPA Service's BIND Name Service for hostname-to-address mapping, you cannot use SAM to remove NFS Services connectivity information about a remote system. The `View/Remove Connectivity Info About a Remote System` form edits only the `/etc/hosts` file; it does not edit an NFS Network Information Service or BIND Name Service database.
- If you were using a specified gateway, other than the default gateway, to reach the remote system you are removing connectivity information about, SAM will automatically remove the special routing information by executing an `/etc/route delete host` command and deleting it from `/etc/netlinkrc`.

To perform this task:

1. Select `View/Remove Connectivity Info About a Remote System`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. If you want to remove the host name from `/etc/hosts`, press **Perform Task**.

4. If you have another task to perform, press Exit Task to go back to the NFS configuration menu. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Prevent This System from Accessing Remote File Systems via NFS (Stop Being an NFS Client)

This task allows you to set up your local system so that it is no longer an NFS client. When you perform this task you are editing the `/etc/netnfsrc` file, killing all `/etc/biod` daemons, and unmounting all NFS- mounted file systems:

1. Select `Prevent This System from Accessing Remote File Systems via NFS`.

   If the NFS configuration menu reads `Allow This System to Access Remote File Systems via NFS`, then your NFS client capabilities are already disabled. Skip to the next task you wish to perform.

2. Follow the instructions in the pop-up window to prevent your system from being an NFS client.

3. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Modify NFS Mount Options

For currently mounted NFS file systems and NFS file systems listed in `/etc/checklist`, this task allows you to modify:

- The name of the local mount directory.
- Whether you want the file system to be mounted now or at reboot.
- Whether or not the file system is write protected.
- The SUID protection.

When you perform the following task, you may be editing the `/etc/checklist` file:

1. Select `Modify NFS Mount Options`.

2. Fill in the form according to its instructions. View the help screens for information about filling in the form.

3. Press Perform Task.

4. If you have another task to perform, press Exit Task to go back to the NFS configuration menu. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Remove (Unmount) an NFS File System

This task allows you to remove the NFS file systems you no longer wish to access from a remote server. When you perform this task, you may be editing the `/etc/checklist` file:

1. Select `Remove (Unmount) an NFS File System`.

2. Fill in the form according to its instructions. View the help screens for information about filling in the form.

3. Press Perform Task.

4. If you have another task to perform, press Exit Task to go back to the NFS configuration menu. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Prevent Remote Systems from Accessing Local File Systems via NFS (Stop Being an NFS Server)

This task prevents remote systems from accessing your local file systems. When you perform this task you are editing the `/etc/netnfsrc` file and killing all `/etc/nfsd` daemons and the `/etc/pcnfsd` daemon:

1. Select `Prevent Remote Systems from Accessing Local File Systems via NFS`.

   If the NFS configuration menu reads `Allow Remote Systems to Access Local File Systems via NFS`, then your system is already disabled as an NFS server. Skip to the next task you wish to perform.

2. Follow the instructions in the pop-up window to prevent your system from being an NFS server.

3. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Caution If any systems have your file systems mounted via NFS, you may cause problems if you stop your system from being an NFS server; their processes that access your file systems will hang if they try either to read from or write to your file systems.

View or Modify Which Systems Can Access Local File Systems

This task allows you to modify a client's NFS access of your local file systems. When you perform this task, you are editing the `/etc/exports` file.

1. Select `View or Modify Which Systems Can Access Local File Systems`.
2. Fill in the form according to its instructions. View the help screens for information about filling in the form.
3. Press Perform Task.
4. If you have another task to perform, press Exit Task to go back to the NFS configuration menu. If you have no further tasks to perform, press Main Menu and Exit SAM to exit SAM.

Note Any changes you make will not affect current NFS mounts by remote systems. However, if these remote systems reboot or unmount the file system(s), your changes may affect them if they try to remount the file system(s).

Prevent systems from accessing local file systems via NFS (without using SAM)

You may need to prevent file access via NFS from either one of the following situations:

- The client to keep local users from accessing NFS-mounted remote file systems.
- The server to keep all clients from accessing local file systems via NFS.

Unmount File Systems from Client

If needed, you can unmount file systems on a client. Unmounting file systems prevents further access to the server's files until you remount the file system. Read the following:

- Executing the `umount` command unmounts file systems mounted either via `mount` or `/etc/checklist`.
- You cannot unmount an open directory or a parent of an open directory (e.g., a directory in which someone is currently operating).
- If operating in an HP-UX cluster environment:
  - If a cnode mounts a remote file system, any cnode in that cluster can unmount the remote file system.
  - If a cnode unmounts a file system, all cnodes in the cluster will have that file system unmounted.

Unmount File Systems on Clients	Action
One NFS file system on a client.	On the client, execute `umount`: `umount` *mount_point_name*
All file systems on one client.	On the client, execute `umount -a`: `umount -a` *Note*: This command unmounts all file systems, not just NFS file systems. If operating in an HP-UX cluster environment, clients should not execute `umount -a`.
All NFS file systems on all clients	On all clients, execute `umount -at`: `umount -at nfs`
All file systems listed in `/etc/mnttab` that were remotely mounted from a specified server	On all clients, execute `umount -h`. `umount -h` *server_name*

Prevent Access to Server File Systems

If needed, you can prevent clients from accessing file systems on the network servers.

Prevent Access to Server File Systems	Action
One NFS file system from a client	1. You have two options for Step 1: - If a netgroup is specified for that file system in `/etc/exports`, remove the host name from the netgroup entry in the server's `/etc/netgroup` file. - If a host name is specified for that file system in `/etc/exports`, remove the host name from the server's `/etc/exports` file. 2. On the client, execute `umount`. `umount` *mount_point_name*
One NFS file system from a netgroup	1. On the server, remove the netgroup name (associated with that file system) from either the `/etc/exports` file or from `/etc/netgroup`. 2. On all members in the netgroup

Prevent Access to Server File Systems	Action
All NFS file systems from all clients	1. On all clients, execute `umount`. `umount` *mount_point_name* 2. On the server, you have two options for Step 2: - Kill the `nfsd(1M)` daemon or daemons (usually four); the system prohibits NFS accesses only until you restart the `nfsd` daemons or you reboot the system. - Edit `/etc/netnfsrc` to change the `NFS_SERVER=` value to zero, and reboot the system. `NFS_SERVER=0`

Update Software

To install a new system release to a server, use the `/etc/update` program to install software. (Refer to the *HP-UX System Administrator's Manual* for detailed instructions.)

The following list includes configuration files loaded during the `/etc/update` process. Some of these files contain example entries to help you configure them correctly:

- `/etc/checklist`
- `/etc/netnfsrc`
- `/etc/inetd.conf`
- `/etc/rpc`
- `/etc/netgroup`
- `/usr/adm/inetd.sec`
- `/etc/netnfsrc2`

Note

If you are mounting file systems, then load *only* those file sets that reside on the local file systems.

For 8.0, `/etc/netnfsrc` has changed significantly. A new `/etc/netnfsrc` file is loaded during the install and update procedure. Standard values are propagated from the old version to the new version. The old version is saved in `/etc/netnfsrc.OLD`. You must copy any customization necessary from the old version.

When using `/etc/update`, the system creates new configuration files in the `/etc/newconfig` directory. These files correspond to the original configuration files which the system leaves in `/etc`.

- Compare each file in `/etc/newconfig` with its existing counterpart in `/etc` to determine if you need to update or replace the file.
- If needed, edit the `/etc/newconfig` files to meet your specific needs.
- Once the `/etc/newconfig` file suits your configuration needs, replace the existing file in `/etc` with the new one in `/etc/newconfig`.
- You may want to save the old configuration file for later reference.

Clock Skew

The NFS client and server clocks may not be synchronized since each workstation keeps its own time. Problems may occur because of these time differences.

If your application depends on the local time or file system timestamps, then it may have to handle clock skew problems if it uses remote files. For example, when giving `utime` a NULL pointer for the times value, the following process occurs:

1. The system sets the access time and modification time according to the client node clock.
2. It then sends these times over to the server, which then changes the inode to reflect the new access and modification times.
3. The server node identifies the change in the inode and thus, modifies the inode's status change time according to its own clock.

The result is a high probability of differing times between the file or directory's access and modification times versus its status change time.

Note HP corrected the clock skew problems that existed with the `ls` command and the source code control command `SCCS`.

If operating in an HP-UX cluster environment, all nodes in the cluster have the same time as the root server's clock. Therefore, clock skew problems exist only if the root server's clock is different from other NFS servers.

EXAMPLE: This example shows how a command could be affected by the clock skew.

Problem Most programs logically assume an existing file could not be created in the future; one example is `ls`. (Note: This example shows how HP corrected this problem.)

The `ls -l` has two basic forms of output, depending on how old the file is.

```
$ date
April 7  15:27:31 PST 1987

$ ls -l file*
-rw-r r 1 root  other   Aug 26  1981   file      (Form One)
-rw-r r 1 root  other   Apr 07  15:26  file2     (Form Two)
```

Form One of `ls` prints the month, day, and year of the last file modification if the file is *more* than six months old. Form Two prints the month, day, hour, and minute of the last file modification if the file is *less* than six months old.

The `ls` command calculates the age of a file by subtracting the modification time of the file from the current time. If the results are greater than six months, the file is "old."

Now assume that the time on the server is three minutes ahead of the local node's time (April 7, 15:30:31). The following commands demonstrate the effect of this clock skew prior to HP's correction of the problem.

```
$ date
April 7  15:27:31 PST 1987
$ touch file3
$ ls -l file*
-rw-r r 1 root  other   0 Aug 26  1981    file
-rw-r r 1 root  other   0 Apr 07  15:26   file2
-rw-r r 1 root  other   0 Apr 07  1987    file3
```

The problem is that the difference of the two times is negative, but the variable in the computation is unsigned. A signed negative number has the same representation (bit pattern) as a very large unsigned number.

local node time = 15:27:31
modification time = local node time plus 180 seconds

local node time	15:27:31
- modification time	- (15:27:31 + 180)
large unsigned number that is greater than six months	- 180 seconds

Problem Correction

HP corrected the problem so that `ls` now prints the month, day, and minute for files between six months old and one hour ahead of time. Other applications may also require such modification.

```
$ date
April 07  15:27:31 PST 1987
$ touch file3
$ ls -l file*
-rw-r rl root other 0 Aug 26 1981   file
-rw-r rl root other 0 Apr 07 15:26  file2
-rw-r rl root other 0 Apr 07 15:30  file3
```

Maintain the NFS Server

NFS servers are described as stateless. This means the NFS server does not know and does not care which clients import its file systems. This gives NFS the advantage of allowing clients to access NFS mounted file systems after the server has recovered from a crash.

To recover as a client, all you need to do is try an NFS remote procedure call. (This is done automatically with hard mounted file systems.)

To recover as a server, just reboot the system (assuming that you have networking and `nfsd` configured to start at boot time).

However, problems occur when servers remain down for extended periods of time and/or file systems are rearranged or modified while the servers are down. In these cases, clients can hang on NFS mounts. In extreme cases, servers' file systems may be corrupted when they become available to NFS clients. There are some precautions you can take to reduce the problems experienced during server maintenance and server crashes. These precautions are covered in the following sections.

Planned Downtime

Before bringing down an NFS server, HP recommends that you unmount all NFS clients. This ensures that no clients hang on a server that is not responding. It also protects the server from possible file system corruption if file systems are modified while the server is down.

However, unmounting NFS clients can be a problem. Unless you restrict access to a small group of clients, it is difficult to determine which clients have file systems imported from a particular server. The `/etc/rmtab` file on the server will give some clues which nodes might be clients, but it is not reliable. The only way to get a complete list of clients is to do an exhaustive search of all possible clients. Do this by logging onto all possible clients, executing the `mount` command, and searching the output for evidence of a file system imported from the server.

If file systems are modified or rearranged on the server while the server capabilities are disabled and not all clients have unmounted the server's file systems, file system corruption may occur when the server resumes servicing NFS requests. To reduce the possibility of file system corruption, run `fsirand` on the server's file systems while they are unmounted locally.

(See *fsirand*(1M) in the *HP-UX Reference*.) This will randomize the inode generation numbers on the file systems thereby minimize the possibility of NFS clients incorrectly modifying these file systems after NFS Services resume on the server.

Unplanned Downtime

When NFS servers crash unexpectedly, NFS clients obviously cannot access file systems imported from the crashed servers. If the clients have hard-mounted NFS file systems, which is the default, client applications that attempt to access those file systems will hang. After rebooting the server, the applications will continue as normal.

Hanging applications can be an annoyance if servers are down for a short period of time. However, hanging applications can become a big problem if servers are down for long periods of time. For this reason, the soft mount option is provided (see "Create an NFS Client Manually"). The soft mount option allows you, as the NFS client system administrator, to determine how many times an NFS request should be transmitted before giving up. This allows applications to continue processing when access to NFS file systems is impossible due to a crashed server.

When NFS servers are not brought down gracefully, there is always the possibility of file system corruption. If the damage is significant, you may need to recreate the file system. If this is the case for a file system that is exported via NFS, it is important to run `fsirand` on the file system after it is recreated and before it is mounted locally. This will add a level of protection on the file system when the NFS client applications attempt to access the file system.

Note Running `fsirand` does not guarantee that a file system corruption will not occur when NFS clients attempt to access the file system. However, it does significantly reduce the possibility.

Remote Execution Facility (REX)

This chapter describes how to configure and execute commands on a remote host using the Remote Execution Facility (REX).

REX consists of:

- The `on` command
- The `rexd` (remote execution daemon)

The `on` command provides the REX user interface on the client. It also communicates with `rexd` to execute commands remotely. `rexd` runs on the server and facilitates the execution of the remote commands.

The functionality of REX is similar to that of remote shell (`remsh`) with two important differences:

- REX executes commands in an environment similar to that of the invoking user. Your environment is simulated by:
  - Copying all of your environment variables to the remote computer.
  - Mounting the file system containing your current working directory on the remote computer via NFS (if it is not already mounted on the remote computer). Your command is then executed on the remote computer in the remote version of your working directory, using your (the invoking user's) environment variables.
- REX allows you to execute interactive commands such as `vi`. In this case, your current `tty` settings (e.g. your current "break" character) are also copied to the remote system.

The on Command

The on command provides the user interface for remote execution of commands. When executing the on command, you specify:

- A host on which to run the remote command.
- The command to run.
- Arguments for the command.

The on command then simulates your current environment on the server by passing your environment variables and information about your current working directory to the remote host. The rexd daemon on the server mounts the file system that contains your current working directory if it is not already mounted on the server. After the environment is simulated, the command executes in the simulated environment on the remote host.

Note Your environment is simulated on the remote host but not completely recreated. Execution of a given command on a remote host will not always produce the same results as executing the command on your local computer. The simulated environment and the environment's limitations are discussed below in "Environment Simulation."

The syntax of the on command is as follows:

```
on [-i | -n] [-d] host [ command [argument ] ....]
```

Host specifies the name of the host on which to execute *command*. There must be an entry for *host* in the local computer's host data base.

Command specifies the command to execute on *host*. If *command* is not specified, on will start a shell on *host*.

You may specify three options (-i, -n, -d). The -i option must be used when invoking interactive commands, the -n option must be used when running commands in the background with job control, and the -d option is used when you wish to receive diagnostic messages.

Use of the -d option with either -i or -n is permitted.

EXAMPLE:

```
on -i -d host
```

or

```
on -n -d host
```

You *cannot* use the -i and -n options at the same time.

The -i Option (Interactive Mode)

The -i option invokes the interactive mode. This option must be specified for all interactive commands (commands which expect to be communicating with a terminal). Examples of interactive commands are vi, csh, and more. If this option is specified with a non-interactive command such as sort, it will be executed as an interactive command, but there may be no difference in behavior.

EXAMPLE:

```
on -i node_7 vi <file>
```

The -n Option (No Input Mode)

The -n option sends the remote program an end-of-file when the program reads from standard input instead of connecting the standard input (stdin) of the on command to the standard input (stdin) of the remote command. The -n option is necessary when running commands in the background with job control.

The -d Option (Debug Mode)

The -d option allows you to receive diagnostic messages during the start up of the on command. The messages may be useful in detecting configuration problems if the on command is failing while connecting to a given host.

Configuration Requirements

The following list details the configuration requirements that must be met for you to execute the on command from node A to node B:

- You must be logged into a user account (other than root) on node A.
- You must have an account on node B, and the UIDs for the accounts on node A and node B must be the same. If this is not the case, one of two things will happen:
  - If the UID associated with the user on node A is not associated with any user on node B, the on command will fail with the error:

    ```
    on hostname: rexd: User id xxxx is not valid.
    ```

  - If your UID on node A is associated with another user on node B, then the command will be executed on node B as the user associated with the UID. (The second case is a serious security limitation. More details are given in the "Security Considerations" section of this chapter).
- The file system that contains your current working directory must be exported in a manner that allows computer B to mount it. Note that the current working directory may be a directory on another remote computer C, which is being accessed via NFS.
- Node B must have rexd configured to execute.

Environment Simulation

As previously mentioned, your environment is simulated on the remote computer, not mirrored. Therefore, certain limitations exist that may cause the execution of a given command to produce different results when executed on the local computer and a remote computer via `on`. These limitations are as follows:

- If the file system is not already mounted on the remote computer, the file system containing your current working directory will be mounted on the remote computer in a subdirectory of `/usr/spool/rexd`. If the file system is already mounted on the remote computer, the mount point is the current mount point for the file system. Therefore, the use of absolute path names can cause problems.

  EXAMPLE:

  User `mjk` on node A is in his home directory (`/users/mark/mjk`) and executes the `on` command to start a shell on a remote system. When the shell is started, the current directory will be `/usr/spool/rexd/rexdAXXXX/users/mark/mjk` (where *A* is a letter and *XXXX* is a 4 digit number). If `mjk` now types the command `cd`, one of two events will occur, depending on the configuration of the file system on the remote computer:

  - If the path `/users/mark/mjk` exists on the remote system, the current directory will be `/users/mark/mjk` on the remote system, which is not equivalent to `/users/mark/mjk` on the local system.
  - If the path `/users/mark/mjk` does not exist on the remote system then executing `cd` will return an error.

  This type of behavior could cause a script that executes `cd` or uses absolute file names to produce different results when executed remotely.

- Another example where the use of absolute path names may occur, without being obvious, is the use of $PATH. Implicit use of $PATH may cause a different version of a command (or a different command) to be executed in the remote case.

- Relative path names will work if they are within the same file system as your current working directory. If a relative path name crosses a file system boundary it will encounter problems similar to those presented by use of absolute path names.

Configuring rexd

Configuring rexd on a system allows the system to act as a server, executing commands for clients that execute an on command. Before configuring rexd to run on a system, you should read the "Security Considerations" section in this chapter.

When rexd is configured, it is started by inetd when a request for remote execution is made by a client. inetd obtains the information it needs to start rexd from the file /etc/inetd.conf. The following entry must be in the file /etc/inetd.conf in order for inetd to start rexd:

```
rpc stream tcp nowait root path 100017 1 rpc.rexd [ options ]
```

Path and *options* are defined as:

path — The path name of the rexd executable in the file system. The rexd shipped with the HP NFS Services product is located in /usr/etc/rpc.rexd.

options — The options that change the behavior of rexd. Each of the possible options is described below:

The -l option

You can log any errors reported by rexd to a file by adding -l log_file at the end of the configuration entry in /etc/inetd.conf, where log_file is the name of the file where errors are logged. If log_file exists, rexd appends messages to the file. If log_file does not exist, rexd creates it. Messages are not logged if the -l option is not specified.

The information logged to the file includes the date and time of the error, the host name, process ID and name of the function generating the error, and the error message. Note that different RPC services can share a single log file since enough information is included to uniquely identify each error.

EXAMPLE:

Thus, the entry in `/etc/inetd.conf` to log errors to the file `/usr/adm/rexd.log` is:

```
rpc stream tcp nowait root /usr/etc/rpc.rexd 100017 1 \
rpc.rexd -l /usr/adm/rexd.log
```

The -m option

Specifying `-m` *mountpoint* changes the default directory containing mount points. This directory is used for mounting client file systems. The following entry in `/etc/inetd.conf` causes client file systems to be mounted as `/client/mnt/rexdAXXXX` instead of `/usr/spool/rexd/rexdAXXXX` (where *A* is a letter and *XXXX* is a 4 digit number):

```
rpc stream tcp nowait root /usr/etc/rpc.rexd 100017 1 rpc.rexd -m /client/mnt
```

The owner, group, and all other users must have read and execute permission for *mountpoint* or an `on` command may fail for a user that does not have the proper permission to *mountpoint*.

The -r option

The `-r` option causes the `rexd` to use stronger security checking than it uses by default (see "Security Considerations"). When started with the `-r` option, `rexd` denies access to a client unless one of the following conditions is met:

- The name of the client is in the `/etc/hosts.equiv` file on the server.
- The user on the server, associated with the UID sent by the client, has an entry in `$HOME/.rhosts` that specifies the client name followed by *one* of the following:
  - White space and an end of line.

    or

  - The user's name and an end of line.

EXAMPLE:

If a user assigned to UID 7 on NODE1 executes the following `on` command:

```
on NODE2 pwd
```

Then user `mjk` (assuming user `mjk` on `NODE2` is assigned UID 7) on `NODE2` must have *one* of the following entries in `$HOME/.rhosts`:

```
NODE1
```

or

```
NODE1 mjk
```

Security Considerations

The design and implementation of REX incorporates several security limitations that you should consider before configuring `rexd`. REX restricts access to a system by use of UIDs. That is, the client (`on`) passes the invoking user's UID to the server (`rexd`) to determine if the invoking user is a valid user. This creates several security limitations:

- If the client and the server do not have the same mapping of user to UIDs, a user on a client may be able to access the server as some other user.
- A malicious user can set the desired UID in the outgoing packets and access the server as any of the server's valid users other than root. An individual with their own workstation can set up a user account with the desired UID.

The impact on system security can be reduced by using the file `/usr/adm/inetd.sec`. The entries in this file specify a set of networks and hosts that are allowed or denied access to a service that is started by `inetd`. For more details on the use of `/usr/adm/inetd.sec` see *inetd.sec*(4) in the *HP-UX Reference*.

The consequences can also be reduced by use of the `-r` option when starting `rexd`. See the previous section, "Configuring rexd," for more details about the `-r` option.

Under normal NFS use, only root is allowed to mount remote file systems. However, when `rexd` is in use, you can mount a file system on the server by executing the following instructions:

1. `cd` to a directory in the file system you wish to mount.
2. Execute the `on` command to start a shell on the computer on which you wish to mount the file system.
3. From another window, shell layer, or system, log into the server and `cd` to a directory in the file system that `rexd` mounted.
4. Switch back to the previous window, shell layer, or system and exit the shell created by the `on` command.

Since another user is busy in the mounted file system, `rexd` will be unable to unmount the file system. Hence, the user has mounted the file system.

Diagnostics

There are two types of error messages discussed in this section. They are:

- `on` command error messages.
- `rexd` error messages.

on Command Error Messages

The following `on` command error messages are sent to `stderr`.

```
on: unknown host <host>
```

The host name <*host*> was not found in the *hosts* database.

```
on: cannot connect to server on <host>
```

The host <*host*> is down, unreachable on the network or not running `rexd`.

```
on: can't find <current_dir>
```

A problem occurred trying to find your current working directory (<*current_dir*>).

```
on: can't locate mount point for <current_dir>
```

A problem occurred trying to determine the mount point of your current working directory (<*current_dir*>).

```
on: standard input (stdin) is not a tty
```

The standard input (`stdin`) of the `on` command with the `-i` option is not a tty.

```
on <server>: rexd: <message>
```

Errors which occur on the server <*server*> are propagated back to the client. These messages are documented in the "DIAGNOSTICS" section of *rexd*(1M) found in the *HP-UX Reference*.

rexd Error Messages

The following is a subset of the messages that may appear in the log file if the `-l` option is used. Some of these messages are also returned to the client.

```
rexd: could not umount <dir>
```

`rexd` was unable to `umount` your current working file system. See *rexd*(1M) in the *HP-UX Reference* for more details.

```
rexd: mountdir (<mountdir>) is not a directory
```

The path name <*mountdir*>, under which temporary mount points are created, is not a directory or does not exist.

```
rexd: <command>: Command not found
```

`rexd` could not find <*command*>.

```
rexd: <command>: Permission denied
```

`rexd` was denied permission to execute <command>.

```
rexd: <command>: Text file busy
```

The executable file is currently open for writing.

```
rexd: <command>: Can't execute
```

rexd was unable to execute <*command*>.

```
rexd: root execution not allowed
```

Root execution is not allowed by rexd.

```
rexd: User id <UID> not valid
```

The UID <*UID*> is not assigned to a user on the server.

```
rexd: User id <UID> denied access
```

rexd was started with the -r option, and the remote execution request did not meet either of the conditions required by the -r option.

```
rexd: <host> is not running mountd
```

The host <*host*> on which the user's current working directory is located is not running mountd. Therefore, rexd is unable to mount the required file system.

```
rexd: not in export list for <file system>
```

The host on which the client's current working directory is located does not have the server on the export list for the file system <*file_system*> containing the client's current working directory. Therefore, rexd is unable to mount the required file system.

6

The Network Lock Manager

Introduction

This chapter explains file and record locking using the Network Lock Manager (`rpc.lockd`) and the Network Status Monitor(`rpc.statd`). It also explains how file locking is used to synchronize access to shared files.

File and record locking allows cooperating processes to synchronize access to shared files. You interface with the networking service by way of the standard `lockf()` system call interface, and rarely require any detailed knowledge of how it works. The operating system maps user calls to `lockf()` and `fcntl()` into Remote Procedure Call (RPC)-based messages to the local lock manager. The fact that the file system may be located on a different node is not really a complication – until a failure occurs.

All computers fail or simply shut down from time-to-time, and in an NFS environment, where multiple computers can have access to the same file at the same time, the process of recovering from a failure is necessarily more complex than in a non-networked environment. Furthermore, locking is inherently stateful. If a server fails, clients with locked files must be able to recover their locks. If a client fails, the locks will be released when the client comes back up. To preserve the overall transparency of NFS, the recovery of lost locks must not require the intervention of the applications themselves. This is accomplished as follows:

- Basic file access operations, such as read and write, use a stateless protocol (the NFS protocol). All interactions between NFS servers and clients are atomic – the server doesn't remember anything about its clients from one interaction to the next. In the case of a server failure, client applications will sleep until the server recovers and NFS operations can complete.
- Stateful services (those that require the server to maintain client information from one transaction to the next) such as the locking service, are not part of NFS. They

are separate services that use the status monitor (see "The Network Status Monitor" section at the end of this chapter) to ensure that their implicit network state information remains consistent with the real state of the network. There are two specific state-related problems involved in providing locking in a network context:

- If the client has failed, the lock can be held forever by the server.
- If the server has failed, it loses its state (including all its lock information) when it recovers.

The Network Lock Manager solves both of these problems by cooperating with the Network Status Monitor to ensure that it is notified of relevant computer failures. The Lock Manager protocol then allows it to recover the lock information it needs when a computer recovers from a failure.

Structure of the Network Locking Service

The following illustration depicts the overall structure of the network locking service.

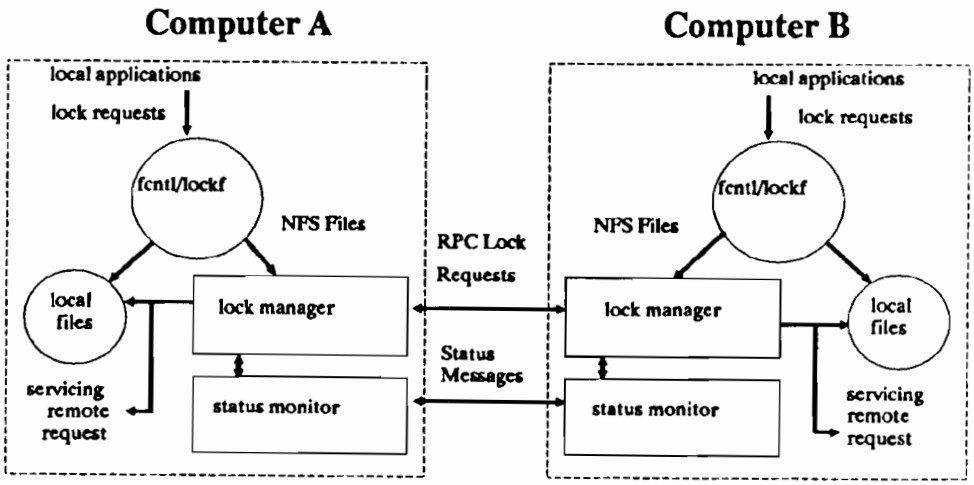

Figure 6-1. Architecture of the Network Locking Service

At each server site, a lock manager process accepts lock requests, made on behalf of client processes by a remote lock manager. The client and server lock managers communicate with RPC calls. When the lock manager receives a remote lock request for a computer that it doesn't hold a lock for, the lock manager registers interest in that computer with the local status monitor. The lock manager then waits for notification from the local status monitor that the computer is up. The local status monitor continues to watch the status of registered computers, and notifies the lock manager when one of them is rebooted (after a failure). If the lock request is for a local file, the lock manager tries to satisfy it, and communicates back to the application along the appropriate RPC path.

If the failure of a client is detected, the server releases the failed client's locks, on the assumption that the client application will request locks again as needed. If the recovery (and, by implication, the failure) of a server is detected, the client lock manager retransmits all lock requests previously granted by the recovered server. This retransmitted information is used by the server to reconstruct its locking state. See the "Locking Protocol" section for more detail.

The locking service, then, is essentially stateless. Or to be more precise, its state information is carefully circumscribed within a pair of system daemons that are set up for automatic application-transparent failure recovery. If a server fails, and thus loses its state, it expects that its clients will be notified of the failure and send it the information that it needs to reconstruct its state. The key in this approach is the status monitor, which the lock manager uses to detect both client and server failures.

Note Failure detection and recovery cannot occur until the system with the failure is rebooted.

Starting the Network Locking Services

The network locking service daemons, `rpc.lockd` and `rpc.statd` are started at boot time by the `/etc/netnfsrc` script. If `/etc/netnfsrc` is configured to make this node an NFS client and/or server, `rpc.lockd` and `rpc.statd` are started by default. If this is not the desired behavior, the `/etc/netnfsrc` script can be modified to produce the desired behavior.

It is important that `rpc.lockd` and `rpc.statd` are running on both the client and server for the file locking protocol to operate correctly.

The Locking Protocol

The lock style implemented by the HP-UX network lock manager supports deadlock detection on a per-server basis only (see the *lockf*(2) and *fcntl*(2) in the *HP-UX Reference* for details).

Despite network lock manager's adherence to the `lockf()` / `fcntl()` semantics, a few subtle points about its behavior need to be mentioned. They are:

- When an NFS client goes down and comes back up, the lock managers on all servers are notified by their status monitors. The server lock managers release locks previously held at the request of the recovering client on the assumption that the client lock manager will request locks again when it wants them. When a server fails, however, the clients wait for the server to come back up. When it does, the server's lock manager gives the client lock managers a grace period to submit lock reclaim requests. During this period the server's lock manager accepts only reclaim requests from remote lock managers. The client status monitors notify their respective lock managers when the server recovers. The default grace period is 50 seconds.

- It is possible that, after a server failure, a client may not be able to recover a lock that it had on a file on that server. This can happen because another process may have accessed the lock before the recovering application process. In this case, the `SIGLOST` signal will be sent to the process (the default action for this signal is to kill the application).

- The local lock manager does not reply to the operating system's lock request until the server lock manager has acknowledged the local lock manager's request. Further, if the lock request is on a server new to the local lock manager, the lock manager registers its interest in that server with the local status monitor and waits for its reply. If either the status monitor or the server's lock manager are unavailable, the reply to a lock request for remote data is delayed until the server becomes available.

- Only advisory mode locking is supported. Enforcement mode is not supported for NFS files.

The Network Status Monitor

The Network Lock Manager relies heavily on the Network Status Monitor to maintain the inherently stateful locking service within the stateless NFS environment. However the status monitor can also be used to support other kinds of stateful network services and applications. Normally, failure recovery is one of the most difficult aspects of network application development, and requires a major design and installation effort. The status monitor simplifies this task.

The status monitor works by providing a general framework for collecting network status information. Implemented as a daemon that runs on all network computers, it uses a simple protocol that allows applications to monitor the status of other computers. Its use improves overall robustness, and avoids situations in which applications running on different computers (or even on the same computer) disagree about the status of a site – a potentially dangerous situation that can lead to inconsistencies in many applications.

Applications that use the status monitor do so by registering the computers they are interested in. The status monitor then tracks the status of those computers, and when one of them fails it notifies the interested applications of the failure, and the applications may take whatever actions are necessary to reestablish a consistent state.

A few advantages of this approach are:

- Only applications that use stateful services must pay the overhead – in time and in size – of dealing with the status monitor.
- The implementation of stateful network applications is simplified, since the status monitor shields application developers from the complexity of the network.

NIS Configuration and Maintenance

The Network Information Service (NIS) is an optional, distributed network lookup service that allows you to administer databases from one node on the network. (NIS was formerly known as Yellow Pages (YP), which is a registered trademark of British Telecommunications.) With NIS you can maintain a single set of user and group IDs for all nodes within a specified set (NIS domain).

The sections included in this chapter are:

- Key terms.
- NIS databases.
- Local and global maps.
- Escape sequences.
- Netgroups.
- Files related to NIS.
- NIS commands.
- NIS configuration.
- Verify NIS.
- NIS maintenance.

Refer to *ypfiles*(4) in the *HP-UX Reference* for a complete explanation of the NIS database and directory structure. If you do not have NIS administrative responsibilities, refer to the *Using NFS Services* manual for general NIS usage information.

Note

If you configured the BIND name server, it will be used instead of NIS for host name and address resolution. However, NIS will still be used for all other information such as passwords. See "Configuring and Maintaining the BIND Name Server" chapter in the *Installing and Administering ARPA Services* manual.

For this chapter only, all references to servers and clients are NIS specific.

Key Terms

Key Terms	Definition
Bind	- Process by which a client locates and directs all requests for data to a specific server. - Process of establishing the address of a socket that allows other sockets to connect to it or to send data to it. Acronym for Berkeley Internet Name Domain. The BIND Name Server is a distributed network lookup service.
Cluster	One or more workstations linked together with a local area network (LAN), but consisting of only one root file system. For more information on cluster concepts, see *Managing Clusters of HP9000 Computers: Sharing the HP-UX Filing System*.
Cluster Auxiliary Server	A cluster client with a disk drive that contains files shared by the other members of the cluster.
Cluster Client	A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.
Cluster Root Server	The only node in an HP-UX cluster that has the root file system directly attached to it.
Escape Sequence (NIS)	Characters used within files to force inclusion and exclusion of data from NIS databases. The escape sequences are as follows: * + (plus) * - (minus) * +@*netgroup_name* * -@*netgroup_name*
Export	To make a file system available to remote nodes via NFS.
File System	An entire unit (disk partition) that has a fixed size.
GID	A value that identifies a group in HP-UX.

Key Terms	Definition
Global (NIS)	A means of access in which the system always reads NIS maps rather than the local ASCII files.
Host	A node that has primary functions other than switching data for the network.
Internet Address	A four-byte quantity that is distinct from a link-level address and is the network address of a computer node. This address identifies both the specific network and the specific host on the network.
Key (NIS)	A string of characters (no imbedded blanks or tabs) that indexes the values within a map so the system can easily retrieve information. For example, in the `passwd.byname` map, the users' login names are the keys and the matching lines from `/etc/passwd` are the values.
Local (NIS)	A means of access in which the system first reads the local ASCII file. If it encounters an escape sequence, it then accesses the NIS databases.
Map (NIS)	A file consisting of logical records; a search key and related value form each record. NIS clients can request the value associated with any key within a map. **NIS map** is synonymous with **NIS database**.
Master Server (NIS)	The node on which one or more NIS maps are constructed from ASCII files. These maps are then copied to the NIS slave servers for the NIS clients to access.
Netgroup	A network-wide group of nodes and users defined in `/etc/netgroup`.
Network Information Service (NIS)	An optional network service composed of databases (maps) and processes that provide NIS clients access to the maps. NIS enables you to administer these databases from one node. NIS may or may not be active; check with your system administrator.

Key Terms	Definition
NIS Client	- A node that requests data or services from NIS servers. - An NIS process that requests other NIS processes to perform operations. *Note* An NIS client can also be configured as any combination of an NIS server, NFS client, or NFS server. An NIS server *must* also be configured as an NIS client.
NIS Database	See "Map (NIS)."
NIS Domain	A logical grouping of NIS maps (databases) stored in one location. NIS domains are specific to NIS and are not associated with other network domains.
NIS Password	The password for a user's login ID that exists in the NIS `passwd` map. The NIS password is the same one as the user password, but is administered through NIS. You do not have to have an NIS password to access the NIS databases.
NIS Server	- A node that provides data (maps) or services to other nodes (NIS clients) on the network using NIS. - An NIS process that performs operations as requested by other NIS processes. *Note*: An NIS server *must* also be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.
Propagate	To copy maps (data) from one NIS server to another.
Slave Server (NIS)	A node that copies NIS maps from the NIS master server and then provides NIS clients access to these maps.
UID	A value that identifies a user in HP-UX.
Value (NIS)	A unit of information stored in NIS maps; each value has a corresponding key (index) so the system can easily retrieve it. For example, in the `passwd.byname` map, the users' login names are the keys and the matching lines from `/etc/passwd` are the values.

NIS Databases

The ypmake script creates the standard NIS databases from the following ASCII files. You can also create additional NIS databases. (Refer to *ypfiles*(4) in the *HP-UX Reference*.)

/etc/group	/etc/passwd
/etc/hosts	/etc/protocols
/etc/netgroup	/etc/rpc
/etc/networks	/etc/services

Other maps may be present, like ethers and mail.aliases, that may be used by other vendors or applications.

Note If the /usr/etc/yp directory is part of a file system that supports only short file names (14 characters maximum), then any maps you create can have only 10 characters. This restriction exists because the makedbm command automatically adds the .dir and .pag suffixes to each map name. Refer to the *System Administration Tasks* manual for more information on short file name file systems.

Local and Global Maps

Clients access the above ASCII files and their corresponding NIS maps in one of two ways, depending on whether the NIS maps are local or global.

- A map is *local* if the system first accesses the local ASCII file. If the file contains an escape sequence, the system then accesses the NIS database.
- A map is *global* if the system accesses only the NIS database (never accesses the local ASCII file).

If a node is not a client, the system accesses only the local ASCII files for information.

NIS Maps	Type	Access
`/etc/group` `/etc/passwd`	Local	If a + (plus) entry exists at the beginning of a line, the system retrieves data from the corresponding NIS map; otherwise, the NIS maps are unused. Occurrences of `+@netgroup_name` and `-@netgroup_name` at the beginning of a line cause the system to reference NIS. (Refer to *group*(4) and *passwd*(4) in the *HP-UX Reference* for complete information regarding these escape sequences.)

`/etc/hosts`
`/etc/netgroup`
`/etc/networks`
`/etc/protocols`
`/etc/rpc`
`/etc/services`
	Global	The system consults only NIS for data. If NIS is not running, it looks at the local ASCII files.
		
		However, if the BIND Name Server is configured, the system will use it for host name and address resolution instead of NIS.

Escape Sequences

Escape sequences are characters used within a file at the beginning of a line to force inclusion and exclusion of data from NIS databases. (Refer to *passwd*(4), *hosts*(4), *netgroup*(4), *host.equiv*(4), and *group*(4) in the *HP-UX Reference*.) The following table shows the escape sequences and their descriptions.

Escape Sequence	Description
+ (plus)	Use + (plus) in `/etc/passwd` and `/etc/group` to retrieve one or more entries from the NIS `passwd` and `group` maps, respectively. The plus designates specific entries to be retrieved from NIS.
- (minus)	Use - (minus) in `/etc/passwd` and `/etc/group` to ignore any subsequent entries with the same name. This process hides the matching names occurring in the NIS `passwd` and `group` maps, respectively. Therefore, it disallows access to particular entries.
+@*netgroup_name*	Use +@*netgroup_name* in `/etc/passwd` to insert the matching entries from the NIS `passwd` map for all members of a network group. *For ARPA Services*: Use +@*netgroup_name* in `/etc/hosts.equiv` and `$HOME/.rhosts` to include a network group's entries from their lists of allowed users.
-@*netgroup_name*	Use -@*netgroup_name* in `/etc/passwd` to disallow the matching entries from the NIS `passwd` map for all members of a network group. For *ARPA Services*: Use -@*netgroup_name* in `/etc/hosts.equiv` and `$HOME/.rhosts` to exclude a network group's entries from their lists of allowed users.

Netgroups

Netgroups are network-wide groups of nodes and users defined in /etc/netgroup on the master server. Use these groups for permission checking during login and remote mount. For *ARPA Services*, you can also use these groups for permission checking during remote login (rlogin) and remote shell execution (remsh).

The master server uses /etc/netgroup to generate three NIS maps in the /usr/etc/yp/domain_name directory: netgroup, netgroup.byuser, and netgroup.byhost. The netgroup map contains basic information found in /etc/netgroup. The other two maps contain more specific information to accelerate the lookup of netgroups given the user or host.

The following table are the programs consulting the NIS netgroup maps and their descriptions. The parenthetical comments refer to sections in the *HP-UX Reference* where you can go for more information.

Program	Description
login(1)	Consults the maps to resolve netgroup names in /etc/passwd.
mountd(1M)	Consults the maps to resolve netgroup names in /etc/exports.
rlogin(1) *remsh*(1)	For *ARPA Services*: Consults the netgroup map if netgroup names are in /etc/hosts.equiv or $HOME/.rhosts.

To limit access to file systems, edit /etc/exports to include the appropriate netgroup names. Then define the netgroup in /etc/netgroup using the following format. (Refer to *exports*(4) and *netgroup*(4) in the *HP-UX Reference*.)

The entry may contain any number of netgroup names:

```
netgroup_name1  netgroup_name2  netgroup_name3  ...
```

But then you must define these netgroups within /etc/netgroup:

```
netgroup_name1  member1  member2  ...
```

You can use the following conventions when editing the /etc/exports file:

- The member*n* is equal to the triple (*host_name, user_name, nis_domain_name*).
- You can assign more than one triple to a netgroup by enclosing each separate set within parentheses (*host_name, user_name, nis_domain_name*).
- Leave any of these three fields empty to signify a wild card (i.e., blank fields match anything). For example, (,,research) matches all hosts and users in the research NIS domain.
- A minus (-) in any of these three fields means *match nothing*. For example, (-,mike,graphs) does not match any hosts, but it does match the user mike in the graphs NIS domain.
- Each host_name must have an entry in /etc/hosts. (See *hosts*(4) in the *HP-UX Reference*.)
- The *nis_domain_name* is the name of the NIS domain to which you currently belong. The commands using /etc/netgroup assume you are not looking for any NIS domain other than the one assigned on your node. (To list your current NIS domain name, execute the domainname command.)

EXAMPLE: The following example is a sample /etc/netgroup file. (Refer to *netgroup*(4) in the *HP-UX Reference* for a complete file format description and a definition of lines and fields.)

```
#
# Engineering: Everyone, but mike, has a node.
# The node 'testing' does not have any users associated with it.
#
engineering hardware software
hardware (mercury,jeff,mickie) (venus,dave,mickie) (testing,-,mickie)
software (earth,carole,mickie) (mars,darren,mickie) (-,mike,mickie)
#
# Marketing: Time-sharing on pluto
#
marketing (pluto,andy,mickie) (pluto,cristina,mickie) (pluto,chuck,mickie)
#
# Others
#
allusers (-,,mickie)
allhosts (,-,mickie)
```

The NIS domain name for all the example netgroups is `mickie`. The following table shows how the users and hosts are classified into netgroups.

Netgroup	Users	Hosts
hardware	jeff, dave	mercury, venus, testing
software	carole, darren, mike	earth, mars
engineering	jeff, dave, carole, darren, mike	mercury, venus, earth, mars, testing
marketing	andy, cristina, chuck	pluto
allusers	every user in the NIS `passwd` map	no hosts
allhosts	no users	all hosts in the NIS `hosts` map

Files Related to NIS

For ARPA Services: The files `/etc/hosts.equiv` and `$HOME/.rhosts` are not in the NIS system; however, they are related to NIS. If these files contain a plus (+) or minus (-) entry with the argument `@netgroup`, the system consults the NIS `netgroup` map for data. (Refer to *netgroup*(4) and *hosts.equiv*(4) in the *HP-UX Reference*.) For example, in `/etc/hosts.equiv` a line consisting of:

```
+@engineering
```

will include all members of engineering as defined in the local file `/etc/netgroup` or in the NIS database. A line consisting only of a plus (+) allows access to all hosts.

The same holds true for `$HOME/.rhosts`. Also, in `$HOME/.rhosts`, a host name followed by a plus (+) means any user coming from that host name will be allowed to access this account through `rlogin` or `remsh`. (See *hosts.equiv*(4) in the *HP-UX Reference*.)

NIS Commands

Refer to the following table for a brief description of all NIS commands. Refer to the "Common Commands" chapter in the "Using NFS Services" manual for a more detailed description of the NIS commands you might want to use on a daily basis (i.e., those NIS commands that do not require superuser access). The parenthetical comments refer to sections in the *HP-UX Reference* where you can go for more information.

NIS Commands	Description
domainname(1)	Use `domainname` to determine or change an NIS domain name.
makedbm(1M)	*Note*: Use this version of `makedbm` only with NIS. A tool for building NIS maps. Use `makedbm` to build or rebuild databases not built by `/usr/etc/yp/ypmake`. Use `makedbm` to disassemble a map so that you can see the key-value pairs comprising it.
ypbind(1M)	Used by each client to determine to which server it should bind. *Note*: This entry exists in the *HP-UX Reference* as *ypserv*(1M); it exists online as *ypbind*(1M).
ypcat(1)	Lists the contents of an NIS map.
ypinit(1M)	On NIS master servers, `ypinit` constructs maps from `/etc` files. On NIS slave servers, `ypinit` copies the initial map versions from the master server.
ypmake(1M)	A script, initially called by `ypinit`, that builds standard NIS maps from ASCII files. These files are usually in `/etc`: `passwd`, `hosts`, `group`, `netgroup`, `networks`, `protocols`, `rpc`, and `services`.

NIS Commands	Description
ypmatch(1)	Prints the value for one or more specified keys in an NIS map.
yppasswd(1)	Changes the password for your current login ID in the NIS `passwd` map. (You do not have to have an NIS password to access the NIS databases.)
yppasswdd(1M)	A server, running only on the master server, that permits users to change their password in the NIS `password` map.
yppoll(1M)	Asks any `ypserv` for the information it holds about a single map.
yppush(1M)	Used by the master server to administer a running NIS system. The `yppush` command causes an NIS map to be copied (using `ypxfr`) from the maps' master server to each slave server in the NIS domain.
ypserv(1M)	Provides access to data stored in NIS maps on servers. If operating in an HP-UX cluster environment, `ypserv` should be running on the root server.
ypset(1M)	Tells the local `ypbind` process to obtain NIS for an NIS domain from a specific server.
ypwhich(1)	Tells you which server a node is currently using or which server is master of a specified map.
ypxfr(1M)	Transfers an NIS map from one slave server to another. Run `ypxfr` *one* of three ways: - `yppush` periodically. - `ypxfr` interactively. - `via cron` periodically.

NIS Configuration

The NIS configuration covered in this section include:

1. Compare `/etc/newconfig` files to existing files.
2. Create an NIS master server.
3. Create an NIS client.
4. Create an NIS slave server.
5. Propagate the NIS maps.
6. Verify NIS.

1. Compare /etc/newconfig Files to Existing Files

When you installed the NFS services software, several new files were copied into the `/etc/newconfig` directory. Perform the following steps to prepare to configure NIS.

1. Compare each `/etc/newconfig` file listed below with its counterpart shown in the following table.

File in /etc/newconfig directory	Counterpart in /usr/etc/yp directory
`ypinit`	`ypinit`
`yp_Makefile`	`Makefile`
`ypmake`	`ypmake`
`ypxfr_1perday`	`ypxfr_1perday`
`ypxfr_1perhour`	`ypxfr_1perhour`
`ypxfr_2perday`	`ypxfr_2perday`

2. If the files are the same, skip to the next section, "2. Create an NIS Master Server."

3. If you have previously customized the files that exist in the /usr/etc/yp directory, or if the files are from an older release of the software, they will differ from files in /etc/newconfig. If there are differences, copy the current files in /usr/etc/yp to a safe location and do *one* of the following:
   - Change the versions in /usr/etc/yp to reflect the differences in the files in /etc/newconfig.
   - Copy the files in /etc/newconfig to /usr/etc/yp. Then re-customize the files in /usr/etc/yp if necessary.

2. Create an NIS Master Server

You must be superuser to create an NIS master server (i.e., to build the NIS master databases). You should also be in a single user state of operation.

An NIS server *must* also be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.

Preparations for Creating an NIS Master Server

Perform the following steps before creating your master server:

1. Ensure /etc files are complete and current: passwd, hosts, group, networks, protocols, rpc, and services.
2. If you know the correct configuration, create the /etc/netgroup file. (See *netgroup*(4) in the *HP-UX Reference*.)

Note The NIS maps store only the first occurrence if:

- A duplicate user name or duplicate user ID exists in /etc/passwd.
- A duplicate internet address or duplicate host name exists in /etc/hosts.

Restricting Access to the Master Server

If you want to restrict access to the master server to a smaller set of users than defined by the complete /etc/passwd file, perform the following steps:

1. Copy the entire /etc/passwd file to a different file (e.g., /etc/passwd.nis).
2. Delete undesired users from the original /etc/passwd file. To prevent all entries in the NIS passwd map from being able to log in, this smaller file *should not* include the following line:

   ```
   +::0:0:::
   ```

3. Edit /usr/etc/yp/ypinit as follows:

   CHANGE: `PWFILE=/etc/passwd`

   TO: `PWFILE=/etc/passwd.nis`

4. Edit /etc/netnfsrc as follows:

   CHANGE: `/usr/etc/rpc.yppasswdd /etc/passwd -m passwd PWFILE=/etc/passwd`

   TO: `/usr/etc/rpc.yppasswdd /etc/passwd.nis -m passwd PWFILE=/etc/passwd.nis`

5. If you have rpc.yppasswdd running, kill and restart it.

   ```
   /usr/etc/rpc.yppasswdd /etc/passwd.nis -m passwd PWFILE=/etc/passwd.nis
   ```

If in the future you need to run ypmake and you have restricted access to the master server as just described, enter the following line:

```
/usr/etc/yp/ypmake passwd PWFILE=/etc/passwd.nis
```

Note For information on C2 Security, refer to the *HP-UX System Security Manual*, *A Beginner's Guide to Using Shells*, and the *HP-UX Beginner's Guide*.

Creating an NIS Master Server

Perform the following steps to create your master server:

1. Set the NIS domain name using the `domainname` command. This NIS domain name must be the same one used for all clients and servers within this NIS domain as shown in the example:

   ```
   domainname nis_domain_name
   ```

2. Execute `ypinit` with the `-m` parameter in one of two ways:

   - If you want to make this node a master server of the domain name that you set in Step 1, enter:

     ```
     /usr/etc/yp/ypinit -m
     ```

   - If you want to make this node a master server of a different domain name than the one you set in Step 1, enter:

     ```
     /usr/etc/yp/ypinit -m [ DOM =      XXX ]
     ```

     XXX represents the domain name for which you are setting this node to be a master server.

3. The system asks whether you want the procedure to quit at the first non-fatal error. Do *one* of the following:

   - Respond `no` or `n` for `ypinit` to continue regardless of the errors. After the procedure finishes, correct all errors that occurred.

   - Respond `yes` or `y` for `ypinit` to quit at the first error. Correct each error as it occurs. This procedure takes longer since you have to correct the errors one by one and run `ypinit` until no more errors occur.

4. The `ypinit` script prompts you for a list of hosts that will become servers.

Starting the NIS Master Server

You should edit /etc/netnfsrc to automatically start the master server at boot time. You can also manually start it now.

Manually Starting NIS Master Server	Automatically Starting NIS Master Server (at Boot Time)

1. If you have not already done so, set the NIS domain name using the domainname command. This NIS domain name must be the same one used for all clients and servers within this NIS domain.
 	
`domainname nis_domain_name`
 	
2. Execute ypserv:
 	
`/usr/etc/ypserv`
 	
Note: If operating in an HP-UX cluster environment, start ypserv only on the node that you wish to make the master server, and start ypbind on every other cnode.
 	
3. Execute ypbind:
 	
`/etc/ypbind`
 	
 
	1. Go into /etc/netnfsrc.
	
	Note: A zero in the NIS_CLIENT, NIS_MASTER_SERVER, or NIS_SLAVE_SERVER field disables the node from working as a client, master server, or slave server respectively.
	
	2. Set NISDOMAIN to the NIS domain name.
	
	`NISDOMAIN=nis_domain_name`
	
	You will need to use this same NIS domain name for all clients and servers within this NIS domain.
	
	3. Set NIS_MASTER_SERVER to a value other than zero. Changing this variable permits users to change their NIS password.
	
	`NIS_MASTER_SERVER=1`
	
	4. Set the NIS_SLAVE_SERVER to zero to disable the node as a slave server.
	
	`NIS_SLAVE_SERVER=0`
	
	5. Set NIS_CLIENT to a value other than zero.
	
	`NIS_CLIENT=1`

3. Create an NIS Client

You must be superuser to create an NIS client.

An NIS client can also be configured as an NFS client, NFS server or both. All NIS servers *must* also be configured as NIS clients. Before creating an NIS client you must:

1. Determine an NIS domain on your network for the client you intend to create.
2. Ensure that a server is available in the NIS domain in which the client will exist (i.e., NIS databases exist and ypserv is running). (Refer to the section "2. Create an NIS Master Server".) If a server is not available in the same NIS domain as the client, users will be unable to log into the client.

Creating an NIS Client

Customize the following files that traditionally store the information. (For suggested modifications, refer to the following section "Altering a Client's Files.")

Note *Do not* abbreviate or eliminate these files if the client is also the master server.

/etc/group	/etc/passwd
/etc/hosts	/etc/protocols
/etc/netgroup	/etc/rpc
/etc/networks	/etc/services

Altering a Client's Files

The following table provides suggestions for altering the client files.

Client File	Suggested Modification
/etc/group	You may want to reduce /etc/group to a single line containing a plus (+) followed by a colon (:) or simply place the line with "+" as the first line of this file. This line forces all translations of group names and group IDs to occur via NIS since group is a local map. +:
/etc/hosts	Ensure /etc/hosts contains an entry for the local host name. The system accesses these entries when NIS is not yet available. After the ypbind process is running, the system never accesses /etc/hosts. EXAMPLE: Sample NIS client's /etc/hosts entry 192.9.1.87 local_host # Byron W. Donnell *Note*: If you configured the BIND name server, it will be used instead of NIS for host name and address resolution. However, NIS will still be used for all other information such as passwords. See "Configuring and Maintaining the BIND Name Server" chapter in the *Installing and Administering ARPA Services* manual.

Client File	Suggested Modification
`/etc/hosts.equiv` (*For ARPA Services*)	The system first accesses `/etc/hosts.equiv` directly. If a `+@netgroup` or `-@netgroup` entry exists, the system accesses the NIS `netgroup` map. *Note*: Using `netgroup` reduces `rlogin` and `remsh` problems that occur because different `/etc/hosts.equiv` files exist on different nodes. For more control over logins, edit `/etc/hosts.equiv` as follows: 1. Enter either a plus (+) or (-) to enable or disable `login`, respectively. 2. Enter the at (@) character. 3. Enter the `netgroup_name` as defined in the global netgroup database. EXAMPLE: *+@netgroup1_name* (trusted) *+@netgroup2_name* (trusted) *-@netgroup3_name* (distrusted)
`$HOME/.rhosts` (*For ARPA Services*)	The system first accesses `$HOME/.rhosts` directly. If a `+@netgroup` or `-@netgroup` entry exists, the system accesses the NIS `netgroup` map. (Refer to the above `/etc/hosts.equiv` example.) Since the superuser's `$HOME/.rhosts` controls remote superuser access to the local node, HP recommends restricted access. To restrict access, either make the list of trusted hosts explicit or use netgroup names.

Client File	Suggested Modification
/etc/passwd	Ensure /etc/passwd contains: - Entries for the root user. - Entries for the primary users. - An escape entry to use NIS. Entries in the local /etc/passwd file mask identical name entries in the NIS passwd maps. Delete all other names and enter +::0:0::: as the last line. This line causes library routines looking for a particular entry to search the NIS database: +::0:0::: EXAMPLES: Sample entries in /etc/passwd +ap::::Dave Hamil:/usr2/ap:/bin/csh The system pulls an entry for ap from the NIS passwd map because of the + (plus) escape character. It obtains the UID, GID, and password from NIS and obtains the comment field, home directory, and default shell from the /etc/passwd entry. If no entry for ap exists in NIS, the system reacts as though no entry for ap exists anywhere: ap::140:100:Mike Donn:/usr2/ap:/bin/csh Since the plus (+) escape character is not present, the system does not access NIS. User ap has no password.

Client File	Suggested Modification
/etc/passwd (*continued*)	EXAMPLES: Sample entries in /etc/passwd +ap: The system obtains all information from the NIS passwd map for user ap. +::0:0::: The system obtains all information from the NIS passwd map for all users not already encountered. If the entry is not found in the NIS database, the search continues in the /etc/passwd file.

Starting the NIS Client

You should edit `/etc/netnfsrc` to automatically start the client at boot time. You can also manually start it now.

Manually Starting NIS Client	Automatically Starting NIS Client (at Boot Time)
1. If you have not already done so, set the NIS domain name using the `domainname` command. This NIS domain name must be the same one used for all clients and servers within this NIS domain. `domainname` *nis_domain_name* 2.Execute `ypbind`. `/etc/ypbind`	1. Go into `/etc/netnfsrc`. 2. Set `NISDOMAIN` to the same NIS domain name used on all clients and servers within this NIS domain. `NISDOMAIN=`*nis_domain_name* 3. Set `NIS_CLIENT` to a value other than zero. `NIS_CLIENT=1` *Note*: A zero in the `NIS_CLIENT` field disables the node from working as an NIS client.

Note If you want the node to be a server also, refer to either the section "2. Create an NIS Master Server" or "4. Create an NIS Slave Server" for complete instructions.

4. Create an NIS Slave Server

You must be superuser to create an NIS slave server.

You may want to create slave servers to improve the reliability of your system. An NIS server *must* be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.

Preparations for Creating an NIS Slave Server

Before creating a slave server, ensure the following:

- The master server exists (see "2. Create an NIS Master Server" section).
- The `ypserv` daemon is running on the master server.

Note If you are operating in an HP-UX environment, HP recommends that only one cnode per cluster should be an NIS server.

Creating an NIS Slave Server

Refer to the following steps to create a slave server:

1. Set the NIS domain name using the `domainname` command. This NIS domain name must be the same one used for all clients and servers within this NIS domain:

   `domainname` *nis_domain_name*

2. Execute `ypinit` with the `-s` parameter in one of two ways:

   - If you want to make this node a slave server of the domain name that you set in Step 1, enter:

     `/usr/etc/yp/ypinit -s` *master_server_name*

- If you want to make this node a slave server of a different domain name than the one you set in Step 1, enter:

  `/usr/etc/yp/ypinit -s` *master_server_name* `[ DOM =` *XXX* `]`

  XXX represents the domain name for which you are setting this node to be a slave server.

3. The system asks whether you want the procedure to quit at the first non-fatal error. Do *one* of the following:
   - Respond `no` or `n` for `ypinit` to continue regardless of the errors. After the procedure finishes, correct all errors that occurred.
   - Respond `yes` or `y` for `ypinit` to quit at the first error. Correct each error as it occurs. This procedure takes longer since you have to correct the errors one by one and run `ypinit` until no more errors occur.
4. Since the slave server is also a client, customize the files which traditionally implement the database. Refer to the previous table "Altering an NIS Client's Files" in the section "3. Create an NIS Client."

Starting the NIS Slave Server

You should edit /etc/netnfsrc to automatically start the slave server at boot time. You can also manually start it now.

Manually Starting NIS Slave Server	Automatically Starting NIS Slave Server(at Boot Time)
1. If you have not already done so, set the NIS domain name using the domainname command. This NIS domain name must be the same one used for all clients and servers within this NIS domain. domainname *nis_domain_name* 2.Execute ypserv. /usr/etc/ypserv *Note*: If operating in an HP-UX cluster environment, start ypserv only on a single node, and start ypbind on every other cnode. 3. Execute ypbind. /etc/ypbind	1. Go into /etc/netnfsrc. *Note*: A zero in the NIS_CLIENT, NIS_MASTER_SERVER, or NIS_SLAVE_SERVER field disables the node from working as a client, master server, or slave server respectively. 2. Set NISDOMAIN to the same NIS domain name used on all clients and servers within this NIS domain. NISDOMAIN=nis_domain_name 3. Set the NIS_MASTER_SERVER to zero to disable the node as a master server. NIS_MASTER_SERVER=0 4. Set NIS_SLAVE_SERVER to a value other than zero. NIS_SLAVE_SERVER=1 5. Set NIS_CLIENT to a value other than zero. NIS_CLIENT=1

5. Propagate NIS Maps

"Propagate NIS maps" means to copy a map from one server to another. Initially, `ypinit` copies the maps when you create slave servers. After the slave servers are initialized, you will use `ypxfr` to transfer updated maps from the master server to the slaves. You can run `ypxfr` three ways:

- Periodically from `cron` on each slave server.
- Periodically by executing `yppush` on the master server.
- Interactively executing `ypxfr` on a slave server.

`crontab`

Maps have different change rates. For example, `protocols.byname` may not change for months, but `passwd.byname` may change several times a day.

Create `crontab` entries to periodically run `ypxfr` at a rate appropriate for each map in the NIS database. The `ypxfr` command will contact the master server and transfer the map only if the master's copy is more recent than the local copy.

To avoid a `crontab` entry for each map, group the maps with approximately the same change characteristics. Place these maps in a shell script you can run via `cron`. Suggested groupings, mnemonically named, are in `/usr/etc/yp`: `ypxfr_1perhour`, `ypxfr_1perday`, and `ypxfr_2perday`. If the rates of change are inappropriate for your needs, either modify or replace these shell scripts.

Execute these shell scripts on each slave server in the NIS domain. Alter the exact time of execution from one server to another to prevent this process from slowing down the master.

EXAMPLE: crontab entries for using these scripts

```
# At 9:00 PM daily, transfer the group, networks, protocols,
# rpc, services, and ypservers maps.

0 21 * * * /usr/etc/yp/ypxfr_1perday

# At 45 minutes past the hour, transfer the passwd maps.

45 * * * * /usr/etc/yp/ypxfr_1perhour

# At 11:30 AM and 11:30 PM daily, transfer the ethers,
# hosts,mail.aliases and netgroup maps.

30 11,23 * * * /usr/etc/yp/ypxfr_2perday
```

You can check and transfer maps with unique change characteristics by explicitly invoking ypxfr from within your crontab file.

EXAMPLE: 25,55 * * * * /usr/etc/yp/ypxfr passwd.byname

yppush

Execute yppush only on the master server to copy a map to each server in the NIS domain (retrieved from the ypservers map).

1. The yppush command sends a "transfer map" request to each of the servers.

2. In turn, ypserv on each server executes ypxfr -C.

3. The ypserv daemon then passes ypxfr the information needed to identify and transfer the map.

EXAMPLE: /usr/etc/yp/yppush passwd.byname

If you wish to run multiple ypxfers at a time and control the timeout value of these transfers, use the -m and -t options. For information about these options, see *yppush*(1M) in the *HP-UX Reference*.

`ypxfr`

Execute `ypxfr` interactively only on the slave servers and only in exceptional situations. For example, execute it when creating a temporary server to make a test environment, or when trying to quickly propagate maps to make a slave server consistent with the other slave servers.

EXAMPLE: `/usr/etc/yp/ypxfr` ***map_name***

If you want the map transferred from a server other than the master server, specify it using the `-h` option with `ypxfr`.

EXAMPLE: `/usr/etc/yp/ypxfr -h server_name passwd.byname`

6. Verify NIS

To verify a client is bound to a server, log into that client and execute `ypwhich`. *One* of the following will occur:

- If the client is bound, the response will be the host name of that server.
- If the client is not bound, you will receive the following message.

  `NIS domain` *`domain_name`* `not bound.`

If you try `ypwhich` several times and continue to receive the `not bound` response, the node is unable to locate a server for that NIS domain on the network. Review your NIS configuration process. If you did not make errors, refer to the "Troubleshooting" chapter.

To verify that NIS is being accessed, log into a client node as a user whose password entry must be served by NIS. If the login does not work, review your NIS configuration process. If you did not make errors, refer to the "Troubleshooting" chapter.

You have now completed configuring NIS. Do *one* of the following:

- If you are configuring NIS for the first time (with NFS Services), and you plan to use the Virtual Home Environment (VHE), you can now skip to the "VHE Configuration and Maintenance" chapter.
- If you do not plan to use VHE, execute `/etc/netnfsrc` to complete the configuration procedure.

NIS Maintenance

To keep NIS running correctly and efficiently, ensure it stays configured to meet your changing needs. Refer to the following sections to help you meet these needs:

- Disable NIS.
- Modify NIS maps.
- Add new NIS servers.
- Add new users to a node.
- Make a different node the NIS master.
- Change NIS password.
- Log files.
- Create non-standard NIS maps.

Disable NIS

You must be super user to disable NIS. If you choose to disable NIS, do the following:

1. Set the NIS domain name to null (no spaces within double quotes).

   ```
   domainname ""
   ```

2. If NIS is currently running, kill the `ypbind` and `ypserv` processes.

3. Edit `/etc/netnfsrc` to change the NIS values:

   a. Change the `NIS_MASTER_SERVER`, `NIS_SLAVE_SERVER`, and `NIS_CLIENT` values to zero:

   ```
   NIS_MASTER_SERVER=0
   NIS_SLAVE_SERVER=0
   NIS_CLIENT=0
   ```

   b. Remove the `NISDOMAIN` variable if one exists:

   ```
   NISDOMAIN=
   ```

4. If the above NIS domain is specified in `/etc/netgroup`, remove the NIS domain name throughout `/etc/netgroup`.

5. Restore any files that you altered for NIS use. For example, you may need to add users back to the `/etc/passwd` file.

6. Reboot the system.

Modify NIS Maps

You must be superuser to modify NIS maps.

Caution Modify maps only on the master server; otherwise, the changes will not be propagated correctly to the slave servers.

You may change most of the standard NIS maps, like `/etc/hosts`, by first editing the ASCII file and then running `ypmake`. Refer to the following "Manual Modifications to NIS Maps" section if you are:

- Adding non-standard maps.
- Editing maps for which no ASCII file exists.
- Changing the set of servers after the system is running.

Whether using `ypmake` in `/usr/etc/yp` or one of the following manual procedures, the goal is the same; a new, well-formed database must reside in the NIS domain directory on the master server. (Refer also to *makedbm*(1M) in the *HP-UX Reference*).

Caution Never modify a map directly; always use `makedbm` to create the map.

Manual Modifications to NIS Maps

You may want to change the following maps manually:

- Non-standard maps (i.e., those that are specific to the applications of a particular vendor or site, but are not part of HP's release).
- Maps that rarely change
- Maps for which no ASCII file exists (e.g., `ypservers` map).

To make a change, do the following:

1. Change to the directory of the maps you want to modify:

   `cd /user/etc/yp/`*`nis_domain_name`*

2. Execute `makedbm -u` to disassemble the map into a form which is modifiable using HP-UX tools:

   a. Redirect the `makedbm -u` output to a temporary file and modify it. Execute `makedbm` using the temporary file as input to create the new versions.

   EXAMPLE:

   ```
   ../makedbm -u mapname > tmpfile
   vi tmpfile  #  (make the required changes)
   ../makedbm tmpfile mapname
   rm tmpfile
   ```

   b. Use a pipe to modify the `makedbm` output which you can then direct as input to `makedbm`. *Note*: You can use this method only if the disassembled map is updated via `awk`, `sed`, or a `cat` append.

   EXAMPLE: Add a new key-value pair to the `map_name` map

   ```
   ( ../makedbm -u map_name; echo  newkey newvalue ) | ../makedbm - map_name
   ```

Examples for Creating Non-Standard NIS Maps

Suppose you want to create a non-standard NIS map. You want it to consist of key-value pairs in which the keys are strings like `a1`, `b1`, `c1`, and `d1`, and the values are `ar`, `br`, `cr`, and `dr`. After creating the map, you notice it is missing `d1` and `dr`.

You could use *one* of two procedures to create the new map: one using an existing ASCII file, the other using standard input.

Example for Existing ASCII File. Assume the following situation:

- An ASCII file exists named `/usr/etc/yp/john_map.asc`.
- The file was created with an editor or shell script on the master server.
- `john_map` is the name of the map you want to recreate.
- `graphs_domain` is the NIS domain subdirectory where the map is located.
- The NIS map was created from this file by entering:

```
cd /usr/etc/yp
./makedbm john_map.asc graphs_domain/john_map
```

Now you notice the map is missing `d1` and `dr`. To correct the error, modify the map by first modifying the ASCII file as follows:

```
cd /usr/etc/yp
<make editorial change to john_map.asc to add the d1 and dr line>
./makedbm john_map.asc graphs_domain/john_map
```

To verify the new map has the changes you made, enter the following command:

```
./makedbm -u graphs_domain/john_map | more
```

Example: Using Standard Input. Assume the following situation:

- `wes_map` is the name of the map you want to create (no ASCII file exists from which the map was built).
- `reports_domain` is the NIS domain subdirectory in which you will create the map.

First, create the NIS map from the keyboard by entering input on the master server as follows:

```
cd /usr/etc/yp
./makedbm - reports_domain/wes_map
a1 ar
b1 br
c1 cr
[CTRL]-[D]
```

To modify the map, use `makedbm` to create a temporary ASCII intermediate file that can be edited:

```
cd /usr/etc/yp
./makedbm -u  reports_domain/wes_map > wes_map.temp
```

Now edit `wes_map.temp` to add the `d1` and `dr` line. Create a new version of the database with the following commands:

```
./makedbm wes_map.temp reports_domain/wes_map
rm wes_map.temp
```

Add or Delete a NIS Server

You must be superuser to add new NIS servers.

If a new slave server is not in the original set, recreate the `ypservers` map on the master server. If needed, rebuild the `hosts` map (refer to *ypmake*(1M) in the *HP-UX Reference*):

1. If the server's address is not in `/etc/hosts`, edit `/etc/hosts` to include the new server's address and then execute `ypmake`:

   ```
   <Edit /etc/hosts>
   /usr/etc/yp/ypmake hosts
   ```

2. Add or delete the host's name to or from the ypservers map in the NIS domain as shown in the following example. Do not delete the master server from the list.

```
cd /usr/etc/yp
./makedbm -u nis_domain_name/ypservers >/tmp/nis_server_list

<Edit /tmp/nis_server_list. Add or delete any slave server>

./makedbm /tmp/nis_server_list nis_domain_name/ypservers
./yppush ypservers
/bin/rm /tmp/nis_server_list
```

3. If you added a slave server, complete the steps in the section "3. Create a NIS Slave Server."

Add New Users to a Node

You must be superuser to add new users to a node.

Refer to the *System Administration Tasks* manual to add new users to a node. The procedure consists of:

1. Editing the master server's /etc/passwd and /etc/group files.
2. Making a home directory.
3. Defining the new user's environment.

Remember to update the NIS passwd and group databases by running /usr/etc/yp/ypmake. If you are using an alternate file to build the NIS password databases, use its full path name instead of /etc/passwd:

```
/usr/etc/yp/ypmake group passwd PWFILE=alternate_passwd_file
```

Make a Different Node the NIS Master

You must be superuser to change the NIS master server to a different node.

1. Copy the following files from your current master server to the node that will be the new master server:
   - /etc/hosts
   - /etc/netgroup
   - /etc/networks
   - /etc/protocols
   - /etc/rpc
   - /etc/services
2. Kill the rpc.yppasswdd process on the current master server.
3. Merge /etc/group and /etc/passwd on the current master server with those on the node that will be the new master server. (If using an alternate password file, you need only copy it.) This merging creates files suitable for building maps for all clients.

   Merging ensures machine-specific password and group entries are kept intact. Either save or delete entries taken from the old master server files. For example, in /etc/passwd save user entries and remove the other node's root entry.
4. If /usr/etc/yp/ypmake, /usr/etc/yp/ypinit, or /usr/etc/yp/Makefile was modified on the old master server to build non-standard maps, copy them and other files from which the non-standard maps are built.
5. On the new master server, complete all steps in the "1. Create a NIS Master Server" section.
6. To prevent starting yppasswdd on the old master server, edit its /etc/netnfsrc file to change the NIS_MASTER_SERVER value to zero:

   ```
   NIS_MASTER_SERVER=0
   ```
7. If the old master server is to be a slave server, complete the steps in the "3. Create a NIS Slave Server" section and the steps in the "2. Create a NIS Client" section.

8. Reboot the new master server.

9. Reboot the old master server.

10. To ensure maps are consistent on all servers, execute `ypinit` on each slave server using the new master server's host name:

```
ypinit -s    new_master_hostname
```

Create or Change NIS Password

The NIS password is the password for a user's login ID that exists in the NIS `passwd` map. The NIS password is used as the user password, but is administered through NIS. Note, you do not have to have a NIS password to access the NIS database.

If you change your password with the `passwd` command, you will change only the entry in your local `/etc/passwd` file if the entry exists. If your password is not in the file, the following error message occurs.

```
Permission denied.
```

If this error occurs, or if you would like to change your password while NIS is in use, execute `yppasswd`.

NIS Password Guidelines

The following list provides the requirements for creating and changing NIS passwords:

- Only the owner or superuser can change a NIS password. The superuser must know the current NIS password to change another user's NIS password.
- Only the first eight characters of the NIS password are significant; the rest are truncated.
- An NIS password must contain at least five characters if it includes a combination of *either one* of the following:
  - Uppercase and lowercase letters.
  - Alpha-numeric characters.

- An NIS password must contain at least four characters if it includes a combination of uppercase letters, lowercase letters, and numeric characters.
- An NIS password must contain at least six characters if it includes only monocase letters.
- You can change an NIS password in the NIS `passwd` map using `yppasswd` only if `rpc.yppasswdd` is running on the master server. (See *yppasswdd*(1M) in the *HP-UX Reference*.)

NIS Password

Refer to the following steps to create or change your NIS password in the NIS `passwd` map:

1. Execute the `yppasswd` command:

   `yppasswd` *user_login_name*

2. The system prompts you for the old NIS password even if one does not exist. If it does exist, enter the old NIS password; otherwise, press [RETURN]. *Note*: The NIS password may be different from the one on your local node.
3. The system prompts you for the new NIS password twice to ensure you enter the correct response. Enter your new NIS password twice, pressing [RETURN] after each entry. The system now updates the master server `passwd` map.

Log Files

Using the `-l` option, you can execute `ypbind`, `ypserv`, and `yppasswdd` so that diagnostic and error messages are written to log files as shown in the following examples:

```
/etc/ypbind -l ypbind_log_file
/usr/etc/ypserv -l ypserv_log_file
/usr/etc/rpc.yppasswdd -l yppasswdd_log_file
```

Preceding each message logged to the file are the date, time, host name, process ID, and daemon name generating the message. Since the messages are uniquely identified by this information, these daemons can share a single log file.

If you execute the daemons without the `-l` option, the following responses occur:.

- The `ypbind` daemon writes its messages directly to the system console, `/dev/console`.
- The `ypserv` daemon writes its messages to the `/usr/etc/yp/ypserv.log` file if it exists when `ypserv` is started.
- The `yppasswdd` daemon provides no messages.

The `ypxfr` command appends transfer information (which map from which server and how many entries it has) to the file `/usr/etc/yp/ypxfr.log` if it exists. The logging occurs only if `ypxfr` is not being run directly by someone at a terminal.

EXAMPLE: Logging occurs if the log file exists and `cron` is running `ypxfr` directly, using a `crontab` entry like the following one:

```
30 * * * * /usr/etc/yp/ypxfr  nis_map
```

All log files could potentially grow without limit until they use up the available file system space. To avoid this occurrence, periodically check the file sizes. One method of preventing this problem is to create a `crontab` entry for each log file as follows:

```
0 1 * * 1,3,5 cat /dev/null > log_file
```

This line truncates `log_file` at 1:00 A.M. every Monday, Wednesday, and Friday.

Create Non-standard NIS Maps

You must be superuser to create and propagate non-standard NIS maps.

The `/usr/etc/yp/ypmake` file supports all of the standard maps shipped by HP. Non-standard maps are those maps which you create that are not originally supported by the `/usr/etc/yp/ypmake` file. To create them:

1. Modify `/usr/etc/yp/ypmake` on the master server so the map can be rebuilt. Modification requirements vary extensively. Generally, though, you need to filter a human-readable ASCII file through HP-UX utilities.

   If the file system in which `/usr/etc/yp` exists supports only short file names

(14 characters maximum), limit the new map name lengths to 10 characters maximum. *Note*: However, the system automatically handles the longer standard NIS map names.

2. If using `Makefile` in `/usr/etc/yp` on the master server to build the maps, modify it so the new map can be rebuilt. (See *ypmake*(1M) in the *HP-UX Reference*.)

3. Modify `/usr/etc/yp/ypinit` on the master server to include the name of your new map in the list of `MASTER_MAPS`. Copy this modified script to all server nodes. This process ensures that any re-initialized or new slave servers will serve the new map.

4. For a client to access the data in the new map, it must exist on each of the servers. Execute the newly modified `ypmake` on the master server to build and copy the map to the current slave servers.

   ```
   /usr/etc/yp/ypmake
   ```

   Slave server support for the propagation of new maps consists of adding `crontab` entries or adding new entries to one of the `ypxfr` shell scripts described in the "Propagate NIS Maps" section.

The following sections cover one example for creating non-standard NIS Maps. The sections of the example include:

- Initial example environment.
- Modify `ypmake`.
- Modify `makefile`.
- Modify `ypinit`.
- Maintain a current access map on each slave server.
- Check the map's contents.

Initial Example Environment

Keep a list of the login names and the host names of all nodes on which each user is allowed to login:

- The information is stored in /usr/etc/access_list.
- The custom NIS map you wish to build from this file is access.

The general form of the ASCII file /usr/etc/access_list is as follows:

```
login_name1 [ host_name1 [ host_name2 ... ] ]

login_name2 [ host_name1 [ host_name2 ... ] ]
.
.
.
login_namen
[ # comments ]
```

- Each user has only one line.
- After each login name are zero or more host names. The user can log into any of these hosts.
- You can use both comments with a # (pound sign) in column one and blank lines.

The following samples could be in /usr/etc/access_list:

```
carole      alpha       catfish       handel
gerbil      catfish

# bigmak is a new hire who has not yet arrived

bigmak
mr_jad      axesys      handel
daveysan    satie       yogurt
chum                    handel
speedy      handel      satie         catfish
fielding    alpha       beta          catfish
```

All of the users except for bigmak can log in on one or more systems.

You may want to use the login name as the key for storing this data in the access map so you can search the map with commands like ypmatch.

```
% ypmatch chum gerbil bigmak carole access
```

In the previous example, ypmatch command would provide an output like the following:

```
chum                    handel
gerbil      catfish
bigmak
carole      alpha       catfish       handel
```

Modify ypmake

Modify /usr/etc/yp/ypmake on the master server as follows.

1. Insert a new function called access() after the services() function:

```
access() {
        grep -v "^[ ]*#" $1 | grep -v "^[ ]*$" | \
        awk 'BEGIN { OFS="\t"; } { print $1, $0 }' | \
        $MAKEDBM - $MAPDIR/access}
```

   This function creates a map that has a key as the first field of each input record, creates a value that is the entire record, and skips over comment lines.

2. Add a new pattern to the case statement that is preceded by "for ARG in $*; do."

   You *must* place this information before the pattern *) in the case statement:

```
access)
        if [ 'expr "$MAPS" : ".* access.*"' -eq 0 ]; then
                MAPS="$MAPS access"
        fi;;
```

3. Add the new map name to the default list of MAPS to build. This addition ensures all maps are built (including the access map) if ypmake is called with no maps specified:

```
MAPS=${MAPS:-'passwd group hosts networks rpc \
    services protocols netgroup access'}
```

4. Add a new pattern to the case statement that is preceded by "for MAP in $MAPS; do":

```
access)       build /usr/etc/access_list access;;
```

Modify Makefile

If using the makefile in /usr/etc/yp on the master server to build the maps, modify it as follows:

1. Insert a new variable called ACCESS after the SERVICES variable:

```
SERVICES = services services.byname
ACCESS = access
```

2. Add the new ACCESS variable to the definition of the ALL_MAPS variable:

```
ALL_MAPS= ${PASSWD} ${GROUP} ${HOSTS} ${NETWORKS} ${RPC} ${SERVICES} \
    ${PROTOCOLS} ${NETGROUP} ${ACCESS}
```

Modify ypinit

1. Modify the /usr/etc/yp/ypinit shell script on the master server to include the new map in list of all maps built on the master server:

```
MASTER_MAPS="group.bygid  group.byname \
    hosts.byaddr hosts.byname netgroup netgroup.byhost \
    netgroup.byuser networks.byaddr networks.byname\
    passwd.byname passwd.byuid protocols.byname \
    protocols.bynumber rpc.bynumber services.byname \
    access"
```

2. Copy this modified script to all current and future NIS servers.

Maintain a Current Access Map on Each Slave Server

1. Execute the newly modified `ypmake` on the master server to build and copy the `access` map to the current slave servers:

```
/usr/etc/yp/ypmake
```

2. On each slave server, modify the appropriate `ypxfr` script to periodically copy the `access` map from the master server:

```
# ypxfr_1perday - Perform daily NIS map check and
#updates
/usr/etc/yp/ypxfr group.bygid
/usr/etc/yp/ypxfr group.byname
.
.
.
/usr/etc/yp/ypxfr access
```

Check the Map's Contents

Execute a few NIS commands to verify the success of your work:

```
% ypwhich -m
services.byname      host1
.
.
.
access          host1
```

This `ypwhich -m` command shows that the server you are bound to now serves the `access` map.

The order of the ypcat listing does not match the order of your file contents:

```
% ypcat access
fielding        alpha           beta            catfish
daveysan        satie           yogurt
speedy          handel          satie           catfish
mr_jad          xesys           handel
gerbil          catfish
carole          alpha           catfish         handel
bigmak
chum            handel
```

The following ypmatch command shows how you can selectively retrieve information from your new access map:

```
% ypmatch speedy daveysan fielding mr_jad access
speedy          handel          satie           catfish
daveysan        satie           yogurt
fielding        alpha           beta            catfish
mr_jad          axesys          handel
```

8

VHE Configuration and Maintenance

Virtual Home Environment (VHE) is an HP-developed service that allows you to configure login environments on remote nodes to mirror the login environment on the users' home nodes. VHE is available to any HP-UX system on a network running the NFS Services product.

You can choose whether to configure and use the service, although when you install NFS Services, VHE is also installed. For an overview of how VHE works, refer to the "NFS Services Overview" chapter.

Note

The Network Information Service (NIS) is not mandatory for using VHE, but this chapter shows how to use VHE assuming NIS is configured and used.

If you do not plan to use NIS, you must have an alternate process for maintaining consistency of the `/etc/passwd` and `/etc/vhe_list` files for all nodes in the VHE group.

Configuration Overview

The following list is an overview of the steps you must complete to configure the nodes on your network with VHE. The steps are described in more detail after the overview list:

1. Prepare for configuring nodes with VHE by obtaining host names for the nodes in your network that will use VHE, installing and configuring NFS Services, and installing and configuring NIS (or instituting an alternate mechanism for maintaining consistent user and group IDs, internet address to host name mappings, password entries and `vhe_list`).
2. Compare VHE files in `/etc/newconfig` directory with existing files in the `/usr/etc/vhe` directory.
3. For each node, decide which file systems are to be mounted and determine the names of mount point directories.
4. Create `/etc/vhe_list` on the NIS master server using the information from step 3.
5. Edit the `/etc/passwd` file on the NIS master server node to contain users' home directories which, in turn, contain the appropriate mount point directories.
6. Distribute the new `/etc/vhe_list` and `/etc/passwd` files by executing `ypmake` on the NIS master server.
7. On each node, edit `/etc/exports`.
8. On each node using VHE, execute `/usr/etc/vhe/vhe_mounter`.
9. Verify that VHE is running correctly.

Note You must be superuser to configure VHE.

1. Complete Preparation Steps

For each node that will use VHE, perform the following steps:

- Obtain a host name.
- Install and configure NFS Services.
- Install and configure NIS (or institute your own mechanism for maintaining consistent host names, group and password entries).

To obtain the host names for the nodes on your network that will use VHE, check your `/etc/hosts` file. If NIS is running, you can use the `ypcat hosts` command to look at the host information. If the BIND Name Server is configured, see the "Configuring and Maintaining the BIND Name Server" in the *Installing and Administering ARPA Services* manual.

To install and configure NFS Services, refer to the "NFS Configuration and Maintenance" chapter.

To install and configure NIS, refer to the "NIS Configuration and Maintenance" chapter. VHE can use NIS for file administration. For VHE to function, it needs all of the nodes using VHE to have a consistent view of the `/etc/passwd` and `/etc/vhe_list` files. NIS provides this; if you are not using NIS, you must ensure consistency by some other method.

The `/etc/vhe_list` file contains a list of all of the nodes that are using NFS to do the same remote mounts. (This is explained in detail in "4. Create /etc/vhe_list.")

NIS maintains single versions of the `/etc/passwd` and `/etc/vhe_list` files on the NIS master server. From the NIS master server, you can add or delete users, change users' home nodes and directories, and add or delete nodes from the VHE group. Once changes are made to `/etc/passwd` and `/etc/vhe_list`, the changes are made in the NIS maps and propagated to the NIS slave servers through the `ypmake` program.

2. Compare /etc/newconfig Files to Existing Files

When you installed the NFS services software, several new files were copied into the `/etc/newconfig` directory. Perform the following steps to prepare to configure VHE.

- Compare each `/etc/newconfig` file listed below with its counterpart shown in the following table.

File in /etc/newconfig directory	Counterpart in /usr/etc/vhe directory
`vhe_mounter`	`vhe_mounter`
`vhe_script`	`vhe_script`

- If the files are the same, skip to the next section, "Determine File Systems and Mount Point Directories."
- If you have previously customized the files that exist in the `/usr/etc/vhe` directory, they will differ from those in `/etc/newconfig`. If there are differences, copy the current files in `/usr/etc/vhe` to a safe location and do *one* of the following:
  - Change the versions in `/usr/etc/vhe` to reflect the differences in the files in `/etc/newconfig`.
  - Copy the files in `/etc/newconfig` to `/usr/etc/vhe`. Then re-customize the newly copied files in `/usr/etc/vhe` if necessary.

3. Determine File Systems and Mount Point Directories

For each node that is using VHE, determine and write down the file systems you want to mount and the directories you want to use as mount points. Use the following conventions when completing this step:

- Begin each mount point pathname with a common path component. (In the examples for this manual, `/vhe` is used.)
- Attach to the above pathname the host name of the machine you plan to mount. For example, for a machine named `vic`, the mount point pathname is `/vhe/vic`. The machine name must match *exactly* the name returned by the `hostname` command (e.g., letters that are in lower case must be typed as lower case and letters that are upper case must be typed as upper case).

- For each file system that will be mounted from each machine to be connected with VHE, attach the file system name to the mount point name. To continue with the above example, if the machine `vic` has two file systems to be mounted: / and `/users`, this would result in the pathnames for the two mount points to be `/vhe/vic/` and `/vhe/vic/users`. In the case of `/vhe/vic/`, you should delete the / at the end of the pathname, resulting in the mount point `/vhe/vic`.

4. Create /etc/vhe_list

The `/etc/vhe_list` file contains a list of all directories that are mount points for your VHE environment. Each node accesses this list for the most current mount point information via NFS mounts. File systems of the remote node are mounted on the appropriate mount point using NFS.

To create the `/etc/vhe_list` file, complete the following items.

- As superuser, edit a file named `vhe_list` in the `/etc` directory of the NIS master server. The `vhe_list` file is installed at the time the NFS product is installed.
- For each mount point on each node create a one-line entry with the following form:

  hostname file_system mount_point [*mount_options*]

  Where:

  - *hostname* is the name of the node whose file system is mounted.
  - *file_system* is the name of the remote file system on the node to be mounted.
  - *mount_point* is the name of the local directory that acts as the mount point for the NFS mount.
  - *mount_options* is an optional field in `vhe_list` that contains options that are passed to the `mount` command. There should be no spaces between items in the *mount_options* field, and the items should be separated by commas. For example, to set the read and write size to 1024 bytes this field would look like:

    ```
    rsize=1024,wsize=1024
    ```

Later, the /usr/etc/vhe/vhe_mounter script uses these fields to perform the appropriate NFS mounts. This script also creates the directories that will be the mount points, so it is not necessary for you to create these directories. If a file exists with the same name as one of the mount point directories, the script produces an error message. In this case, you need to either change the name of the existing file or change the name of the mount point directory.

If you are not using NIS, after you create the /etc/vhe_list file you need to distribute the /etc/vhe_list file to all the nodes in the VHE group.

Example: Simple Configuration with Single File System per Node

In the simplest case, each node has only one file system which is the root file system. Every node needs to have a set of directories for all members of the group. For example, consider a group consisting of the nodes A, B, C and D. A list of mount points for this group is /vhe/A, /vhe/B, /vhe/C and /vhe/D. Now taking these two lists, an /etc/vhe_list file with the following contents is created:

```
A  /  /vhe/A
B  /  /vhe/B
C  /  /vhe/C
D  /  /vhe/D
```

Example: Node with Multiple File Systems

Note If you do not have multiple file systems on each node, you can go to "5. Update /etc/passwd."

Doing mounts of several file systems from one node requires some care in creating the /etc/vhe_list file. For example, if /usr is a separate file system on node C, and you execute the following on node A:

```
mount C:/  /vhe/C
```

An ls of /vhe/C/usr on node A shows it as an empty directory because NFS allows access to separate file systems only if they are explicitly mounted.

This directory can be used to do a `mount` of the `/usr` file system of node C by executing the following on node A:

```
mount C:/usr  /vhe/C/usr
```

Now an `ls` of `/vhe/C/usr` on node A shows the contents of the `/usr` file system on node C.

The example group is changed to show this complication with additional file systems:

```
A  1 file system under  "/"
B  2 file systems one under  "/"
   and one under  "/users"
C  2 file systems one under  "/"
   and one under  "/usr"
D  1 file system under  "/"
```

When a node has multiple file systems, you may choose to have all the file systems mounted (as with C) or to have only some of the file systems mounted (as with B). When `/usr/etc/vhe/vhe_mounter` is run, the mount point directories are created, if necessary, and the NFS mounts are made.

Using the rules outlined in "4. Create /etc/vhe_list," for the above group of nodes, you would create the following `/etc/vhe_list` file:

```
A  /         /vhe/A
B  /users    /vhe/B/users
C  /         /vhe/C
C  /usr      /vhe/C/usr
D  /         /vhe/D
```

5. Update /etc/passwd

Update the `/etc/passwd` file on the NIS master server to force home directory access through the mount points. The entries in `/etc/passwd` should have the following form:

login_name`:`*encrypted_password*`:UID:GID:comment:/vhe/`*hostname*`/`*home_dir*`:shell`

Note If you are not using NIS, after updating the `/etc/passwd` file, you must distribute the changes to all nodes in the VHE group.

Example: /etc/passwd file entries before and after the VHE configuration

In this example, the first user 's home directory is on node A; the second user's home directory is on node B; and the third user's home directory is on node C. All of the `/users` directories are in the root file systems on their respective nodes.

Before VHE configuration:

```
andy::117:100:andy:/users/andy:/bin/csh
speedy::118:100:darren:/users/speedy:/bin/ksh
chum::119:200:Cris:/users/chum:/bin/sh
```

After VHE configuration:

```
andy::117:100:andy:/vhe/A/users/andy:/bin/csh
speedy::118:100:darren:/vhe/B/users/speedy:/bin/ksh
chum::119:200:Cris:/vhe/C/users/chum:/bin/sh
```

Example: Nodes with Multiple File Systems

Nodes with multiple file systems do not change how the home directories are updated for VHE. For example, consider the following two entries in `/etc/passwd`. Fielding's home node is node B, which has two file systems; Jeff's home node is node C, which has two file systems. The nodes are from the example shown above.

Before VHE configuration:

```
Fielding::120:200:fielding:/users/fm:/bin/csh
Jeff::121:100:Jeff:/users/jbrl:bin/csh
```

After VHE configuration:

```
Fielding::120:200:fielding:/vhe/B/users/fm:/bin/csh
Jeff::121:100:Jeff:/vhe/C/users/jbrl:/bin/csh
```

For node B, /users is its own file system and is mounted on the directory /vhe/B/users. This causes no change in the naming convention for the home directory. For node C, /users is on the root (/) file system. Node C also has another file system: /usr. If Jeff wants to be able to change the default pathname to his mail file from /usr/mail/jbrl to /vhe/C/usr/mail/jbrl (to read mail via VHE), the /usr file system must be mounted on /vhe/C/usr.

6. Update /etc/exports

On each node that needs to export file systems, edit the /etc/exports file to reflect all of the file systems that are available for NFS mounting from each node. Details on this can be found in the "NFS Configuration and Maintenance" chapter.

7. Distribute /etc/vhe_list and /etc/passwd

To distribute the /etc/vhe_list and /etc/passwd files (i.e., make them accessible to all the nodes using NIS that are part of the same NIS domain), execute the following command on the NIS master server.

```
/usr/etc/yp/ypmake
```

This builds the NIS maps and propagates the maps to the NIS slave servers.

8. Execute /usr/etc/vhe/vhe_mounter

Note The /usr/etc/vhe/vhe_mounter script should be run when all nodes in the VHE group are powered up and ready for NFS mounting. If they are not ready for NFS mounting, then error messages are printed. These are not fatal errors; to recover from them you should retry vhe_mounter when the nodes are available for mounting.

The /usr/etc/vhe/vhe_mounter script uses the information in /etc/vhe_list to create the appropriate mount point directories on each node. When vhe_mounter notices that it is about to make a directory with the same name as the node from which vhe_mounter is executed, it makes a symbolic link with the same pathname and links it to the node's root directory. When the vhe_mounter process completes running on each node, the proper mount points and symbolic links are created for each node.

The /usr/etc/vhe/vhe_mounter script also does NFS mounts using the appropriate directories to the remote machines on each node. When the mounts are complete, a node is ready for VHE.

To execute /usr/etc/vhe/vhe_mounter for each node separately, execute the following script on each node:

```
/usr/etc/vhe/vhe_mounter
```

To run /usr/etc/vhe/vhe_mounter for all nodes using VHE from a single node, execute the following as a batch file.

```
for i in ' ypcat vhe_list | awk '{ print $1 } ' | sort -u '
do
  remsh $i /usr/etc/vhe/vhe_mounter
done
```

Note For this script to execute correctly, all nodes must be running ARPA/Berkeley Services with superuser capability allowed between the nodes when using remsh.

Example: This example shows the mount points and symbolic links resulting from the following /etc/vhe_list file:

```
A  /    /vhe/A
B  /    /vhe/B
C  /    /vhe/C
D  /    /vhe/D
```

The listing below shows the mount points and symbolic links for each node after the `/usr/etc/vhe/vhe_mounter` script completes running on each node (symlink = / denotes a symbolic link to the root (/) directory):

Node	/vhe/A	/vhe/B	/vhe/C	/vhe/D
A	symlink = /	Directory	Directory	Directory
B	Directory	symlink = /	Directory	Directory
C	Directory	Directory	symlink = /	Directory
D	Directory	Directory	Directory	symlink = /

9. Verify that VHE is Correctly Configured

To check if VHE is configured correctly, pick a login name that had a mount point added to its home directory. After `/usr/etc/vhe/vhe_mounter` has been run on each node, go to each node and log in using that selected login name (with the appropriate password). If VHE is correctly configured, the logins are successfully completed, and you are always placed in the execution environment associated with the selected login name.

Note

You have now completed configuring the VHE service. The following sections describe advanced usage or set-up problems you may encounter when using VHE.

If you are configuring VHE as part of the NFS Services configuration, execute `/etc/netnfsrc` to complete the configuration procedure.

Configuration Refinements

The configuration procedure presented in the previous sections addresses most configuration cases. However, you may wish to refine your VHE configuration. This section explains how to refine your VHE configuration to allow NFS mounts to be done in the background.

NFS mounts in the Background

You can alter the `/usr/etc/vhe/vhe_mounter` script to allow mounts to be done in the background. This eases the situation where all nodes are not ready to respond when a node tries to mount them. To mount nodes in the background, you need to edit the `/usr/etc/vhe/vhe_mounter` script.

The `vhe_mounter` file has a shell variable called BACKGROUND_MOUNT whose initial value is set to 0. To allow nodes to be mounted in the background:

- Use an editor to set the value to something other than 0.
- Save the file and execute the `/usr/etc/vhe/vhe_mounter` script.

These changes cause NFS mounts to occur in the background. If the mounts are not successful on the first try, the NFS mounts continue to execute in the background.

Note Because each mount executes as a separate process until it completes or until the `retries` option for the NFS mount is exceeded, there may be a problem if there are many nodes (more than 30) in the VHE group.

VHE Maintenance

To keep VHE running correctly and efficiently, refer to the following sections.

Unmounting file systems

If needed, you can unmount all of the remotely mounted file systems. The easiest method of doing this is to execute the following:

```
umount -a -t nfs
```

This command can only be used when there are no VHE users logged on. If VHE is currently being used, the mount point directories will be busy and `umount` will not unmount a directory that is busy.

Just as having multiple file systems available for remote mounting required mounting to be done in a specific order, unmounting file systems must be done in the proper order. The order is just the reverse from the order that the mounts were done. The `umount` command with the "-a -t" options does this automatically.

For example:

```
mount A:/ /vhe/A
mount A:/usr /vhe/A/usr

umount /vhe/A/usr
umount /vhe/A
```

Adding or Deleting VHE Nodes

You may need to add or delete nodes from the VHE configuration. To do this, you need to perform the following steps:

1. Update the `/etc/vhe_list` on the NIS master server by either removing file systems that are no longer available (if a node is being deleted) or adding file systems that you want to become available (if a node is being added). Refer to the section in this chapter called "3. Create /etc/vhe_list" for more information about how to do this.

2. Edit the `/etc/passwd` file to show the addition of mount points to the home directory pathname. Refer to the section in this chapter called "5. Update /etc/passwd" for more information on how to do this. If you are removing file systems, you need to edit this file to delete mount points from the home directory pathname.

3. To distribute the `/etc/vhe_list` and `/etc/passwd` files to the NIS servers, execute the following command on the NIS master server:

   ```
   /usr/etc/yp/ypmake
   ```

4. Then execute the following:

   ```
   /usr/etc/vhe/vhe_mounter
   ```

   The script uses the information found in `/etc/vhe_list` to decide which new file systems to mount. The `/usr/etc/vhe/vhe_mounter` script does not attempt to unmount a node deleted from the group. `vhe_mounter` needs to be executed on all of the nodes in the group for all of the nodes to be updated.

Advanced Usage

Adding altlogin and mounter Logins

The two logins of altlogin and mounter can be added to /etc/passwd by the superuser. This allows the user to:

- Log in using the mounter ID to complete NFS mounts to a node, if for some reason a node was not mounted when vhe_mounter was executed.
- Log in using altlogin to access the node where they currently are. This is useful if their home node is down.

These logins are similar to who and date because they execute a program. Mounter executes vhe_u_mnt, and altlogin executes vhe_altlog as follows:

- The vhe_u_mnt program executed by the mounter login only attempts to mount a file system of a node that is found in the /etc/vhe_list file. This prevents users from performing mounts to arbitrary nodes. Users can only perform mounts that could have been done by /usr/etc/vhe/vhe_mounter. If the node name entered at the prompt is not found in /etc/vhe_list, then an error message is printed and the mounts are not completed.
- The vhe_altlog program executed by altlogin prompts for a login ID and then attempts to do a su using the provided login ID. The user is then prompted for a password by su. If the proper password is given, the user is logged in with the home directory of /tmp. (If a proper password is not given, the user is not allowed access to the system.) Once logged in, none of the user's execution environment is available, but he or she can use the system.

To make these logins valid, you need to add them to the /etc/passwd file. Do this by adding an entry for each login to the /etc/passwd file. These entries should be similar to the following:

```
mounter::6:1::/:/usr/etc/vhe/vhe_u_mnt
altlogin::6:1::/tmp:/usr/etc/vhe/vhe_altlog
```

The values shown in the above lines in UID, GID and home directory can be replaced with other values. Also note no password is provided in the above lines, but passwords can be

entered if desired. If passwords are entered, tell the users allowed to use those logins what the associated passwords are because they *must* provide them when logging in.

Mounter Example

In this example, dave attempts to log in from node B when his home node, node A, is not mounted on node B. The following sequence would occur:

```
login: dave
Password:
Unable to change directory to  /vhe/A/users/dave

login: mounter
Password:
Enter the name of the node to mount: A

login: dave
Password:
<Dave  gets  logged   in>
```

Altlogin Example

This section shows an example of using altlogin. Julia is currently working at node B. Her home node A is not up, but Julia can gain access to node B in the following way:

```
login: altlogin
Enter your login name: Julia
Password:
%
```

Julia is now logged in at node B.

$HOME

If you are writing scripts that make reference to files in a home directory, those file names should be prefixed with $HOME (for sh or ksh). For csh, file names should be prefixed with a ~ character. This allows a file to be accessed in a consistent manner even if the home directory pathname changes.

$ROOT

To make a distinction between system files (like the password file) for the local and the home nodes, the following can be added to the `.profile` or `.login` file (`home_node` should be replaced with the name of the node):

```
ROOT=/vhe/home_node
export ROOT
```

This allows easier access to system files on a user's home node. For example, instead of typing:

```
more /vhe/home_node/etc/passwd
```

The user types:

```
more $ROOT/etc/passwd
```

Alternate Mount Points

The mount examples in this chapter are prefixed with `/vhe`. In addition to `/vhe` mount points, there may be other file systems users in a VHE group want to regularly access.

For example, in a given VHE group, node A has file system `/Design`. To have a consistent view of this file system among all users in the VHE group, the `/Design` file system can be mounted on a pathname `/Design`. To do this, the following line would be added to the `vhe_list` file:

```
A /Design /Design
```

Using VHE for Mail

To extend VHE to handle mail tasks:

- Change your default mailbox pathname to have a mount point added to the beginning of it (just as the home directories are changed in `/etc/passwd`).
- Specify the above pathname as the file to be used by the mail handler of your choice. If that mail file is on a separate file system, it must also be mounted to be available.

For example, if user `fm`'s home node is A, this shows how the mailx program can be invoked to read mail over NFS:

```
mailx -f /vhe/A/usr/mail/fm
```

In this example, if `/usr` was a separate file system on A, then the following would be added to `/etc/vhe_list`:

```
A /usr /vhe/A/usr
```

9

Troubleshooting

If a node on the network is not operating correctly, use this chapter to identify and correct the problem. Most problems occur when:

- Installing the network.
- Changing the network (e.g., adding a node or extending the coaxial cable).
- Another system on the LAN fails.

Before troubleshooting the problem, get or create your network map as described in the *Installing and Administering LAN* manual. Use this map when checking configuration and network layout information. Remember to update it any time you make a change to the network.

Note All references to servers and clients apply to NFS servers and clients unless otherwise specified.

Key Terms

Term	Definition
Client	A node that requests data or services from other nodes (servers). A process that requests other processes to perform operations. *Note*: An NFS client can also be configured as any combination of an REX server, NIS client, or NIS server. (An NIS server must also be configured as an NIS client.)
Cluster	One or more workstations linked together with a local area network (LAN), but consisting of only one root file system. For more information on cluster concepts, see *Managing Clusters of HP9000 Computers: Sharing the HP-UX Filing System*.
Cluster Auxiliary Server	A cluster client with a disk drive that contains files shared by the other members of the cluster.
Cluster Client	A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have its root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.
Cluster Node (Cnode)	Any node operating in an HP-UX cluster environment, including cluster clients and cluster servers.
Cluster Root Server	The only node in an HP-UX cluster that has the root file system directly attached to it.
Daemon	Background programs that are always running, waiting for a request to perform a task.
Export	To make a file system available to remote nodes via NFS.
Hard Mount	A mount that causes NFS to retry a remote file system request until it succeeds, you interrupt it (default option), or you reboot the system.
Heterogeneous Cluster	A diskless cluster with more than one type of computer attached.

Term	Definition
Homogenenous Cluster	A diskless cluster composed of nodes of only one type of computer architecture (e.g., HP 9000 Series 300)
Host	A node that has primary functions other than switching data for the network.
Map (NIS)	A file consisting of logical records; a search key and related value form each record. NIS clients can request the value associated with any key within a map. **NIS map** is synonymous with **NIS database.**
Master Server (NIS)	The node on which one or more NIS maps are constructed from ASCII files. These maps are then copied to the NIS slave servers for the NIS clients to access.
Mount	To obtain access to a remote or local file system or directory (import).
Mount Point	The name of the directory on which a file system is mounted.
Netgroup	A network-wide group of nodes and users defined in `/etc/netgroup`.
Network Information Service (NIS)	An optional network service composed of databases (maps) and processes that provide NIS clients access to the maps. NIS enables you to administer these databases from one node. NIS may or may not be active; check with your system administrator.
NIS Client	A node that requests data or services from NIS servers. An NIS process that requests other NIS processes to perform operations. *Note*: An NIS client can also be configured as any combination of an NIS server, NFS client, or NFS server. (An NIS server must also be configured as an NIS client.)
NIS Database	See "Map (NIS)."
NIS Domain	A logical grouping of NIS maps (databases) stored in one location. NIS domains are specific to NIS and are not associated with other network domains.

Term	Definition
NIS Map	See "Map (NIS)."
NIS Password	The password for a user's login ID that exists in the NIS `passwd` map. The NIS password is the same one as the user password, but is administered through NIS. You do not have to have a password to access the NIS databases.
NIS Server	A node that provides data (maps) or services to other nodes (NIS clients) on the network using NIS. An NIS process that performs operations as requested by other NIS processes. *Note*: An NIS server must also be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.
Node	A computer system that is attached to or is part of a computer network.
Server	A node that provides data or services to other nodes(clients) on the network. A process that performs operations as requested by other processes. *Note*: An NFS server can also be configured as any combination of an NFS client, NIS client, or NIS server. (An NIS server must also be configured as an NIS client.)
Slave Server (NIS)	A node that copies NIS maps from the NIS master server and then provides NIS clients access to these maps.
Soft Mount	An optional mount that causes access to remote file systems to abort requests after one NFS attempt.

Troubleshooting References

Troubleshooting the NFS Services primarily concerns the areas: power up and connectivity, NFS Services, NIS, VHE, and REX. This chapter only addresses NFS, NIS, VHE, and REX problems. Link diagnostics and troubleshooting are in the *Installing and Administering LAN* manual.

If your system is having problems communicating with or through a non-HP system, refer also to the appropriate user and system administration documentation for that system.

Power Up and Connectivity Testing

Refer to the following documentation if your system cannot communicate with other systems on the network:

- *LAN Interface Controller (LANIC) Installation and Reference Manual.*
- *HP Repeater Installation Manual* (only if you are using a HP 92223A repeater).
- *Installing and Administering LAN/9000.*
- *Installing and Updating HP-UX.*
- *HP-UX Reference.*
- *System Administration Tasks.*
- *LAN Cable and Accessories Manual.*
- *Installing and Administering Network Services.*
- *Installing and Administering ARPA Services.*

Troubleshooting Sections

Refer to the "Troubleshooting NFS" section or the *HP-UX Reference* if you cannot mount a remote file system, access a remotely mounted file system, or experience other problems with the NFS service.

Refer to the "Troubleshooting NIS" section or the *HP-UX Reference* if you configured the system to use the Network Information Service, but cannot access files serviced by it.

Refer to the "Troubleshooting VHE" section if you configured the system to use VHE, but it doesn't function as described in the "VHE Configuration and Maintenance" chapter.

Refer to the "Troubleshooting REX" section if you configured the system to use REX, but it doesn't function as described in the "Remote Execution Facility (REX)" chapter.

Guidelines

Troubleshooting is an elimination process that narrows a problem. If a process worked before but does not work now, first consider what has changed. For example, have you moved hardware or modified configuration files?

Start with the minimum number of variables, then gradually and selectively add other variables such as the following:

- If you cannot communicate with one system, try another one. If the second system works, the problem may be with the first remote system and not your system.
- If one system cannot communicate with yours, try another one. If neither system can communicate with yours but they can communicate with each other, the problem may be with your system.
- If one service does not work, try another one. The problem may be with a particular service to a particular system and not a problem with the system itself.

Common Network Problems

Network problems generally occur under the following circumstances:

- File permissions on the client or server restrict the operation.
- Network services on the client or server are misconfigured or malfunctioning.
- Network LAN software or hardware is misconfigured or malfunctioning.

Initial Troubleshooting

You should first check the following situations to ensure they are not the cause. If they are not, refer to the flowcharts in this chapter.

Configuration

1. Is your host running HP-UX 6.0 or later for the Series 300/400 and HP-UX 2.0 or later for the Series 600/700/800? For File Locking and REX, your host must be running HP-UX 6.5 or later for the Series 300/400 or HP-UX 7.0 or later for the Series 600/700/800. Execute `uname -a` or `uname -r` to check the HP-UX version number.

2. Does your system have the recommended 256K additional memory for networking software?

3. Is your HP 9000 a supported configuration? If you are unsure, contact your HP support representative.

4. Does the error occur on a node other than a Series 300/400 or Series 600/700/800? If so, refer to the appropriate system documentation.

Hardware

The *Installing and Administering LAN/9000* documentation contains details about troubleshooting hardware problems.

1. Are all connections along the network cabling tight?

2. Is each cable segment less than 500 meters for ThickLAN and less than 100 meters for ThinLAN?

3. Are there no more than two repeaters between you and the node with which you want to communicate?

4. Are you mixing Ethernet[1] hardware with IEEE 802.3[2] hardware? This is not an acceptable combination since they do not have the same electrical characteristics.

5. Is there a 50 ohm terminator at the end of each cable?

6. Is the MAU tapped correctly into the cable?

---

(1) Ethernet is a local area network system developed by Digital Equipment Corporation, Intel Corporation, and Xerox Corporation.

(2) IEEE 802.3 is a networking standard that is accepted by the Institute of Electrical and Electronic Engineers.

7. Is the cable grounded in only one place?
8. Is the AUI solidly connected to the interface card?
9. Is the host hardware working correctly?

Network Communication

1. Is the remote node HP certified? If you are unsure, contact your HP support representative.
2. Can any other two nodes on the network communicate? If not, the problem may be global. Refer to the *LAN Cable and Accessories Installation Manual* and *Installing and Administering LAN/9000* documentation.
3. Have you performed the corrective action supplied with the error message you received? Consult the appropriate entry in the *HP-UX Reference*.
4. If using gateways, do both hosts have routing information to each other? Refer to *route*(1M) in the *HP-UX Reference*.
5. If operating in an HP-UX cluster environment and trying to mount an NFS file system, ensure you are using the cluster server's host name (on which the file system is mounted) as the node specified in the `mount` command. This will be either the cluster root server or the cluster auxiliary server.
6. If operating in an HP-UX cluster environment and having link problems, cnodes will not be able to boot. Since link diagnostics reside on the root disk, first test the Link from the root server. (Refer to *Installing and Administering LAN/9000* documentation.)

NIS and NFS Services

1. Is the client system trying to perform tasks as superuser on the remote system? Executing `setuid` root programs cannot access files or directories unless the permission `other` allows it.

2. Was network communication established between the client and server using the procedures outlined in the *System Administration Tasks* manual, and in the "NFS Configuration and Maintenance" and "NIS Configuration and Maintenance" chapters of this manual?

3. Is the problem associated with remote file locking? The *lockf*(2) call fails when attempting to lock a remote file. Prior to HP-UX release 6.5 for the Series 300/400 and HP-UX release 7.0 for the Series 600/700/800, NFS Services *did not* support file locking on remote file systems.

4. Is the problem associated with attempts to access remote device files? Prior to release 6.5 for the Series 300/400 and HP-UX release 7.0 for the Series 600/700/800, HP-UX *did not* support remote access to device files.

5. Does the `inetd` security file (`/usr/adm/inetd.sec`) on the remote system limit access to the remote system for the RPC service you are trying to access?

6. Is the file system listed in the server's `/etc/exports`?

7. Does `/etc/exports` restrict file system access to a specific netgroup or host?

   a. The `/etc/netgroup` file must list the netgroup if it is specified in `/etc/exports`.

   b. The `/etc/hosts` file must contain the host if it is specified either in `/etc/exports` or in `/etc/netgroup`.

8. Is the file system or directory mounted? To check, execute the `mount` command.

9. If the file system is suppose to be automatically mounted, is it listed in `/etc/checklist`?

10. If programs accessing remote files hang, is the NFS or NIS server down?

11. Is data on remote nodes corrupted? Ensure only one system is writing to the file at a time; NFS allows more than one client to write to a file simultaneously.

Remote Execution (REX)

1. Is the server configured to run `rexd`? The server must have an entry in `/etc/inetd.conf` in order to run `rexd` (see *rexd*(1M) in the *HP-UX Reference*).

2. Was network communication established between the REX client and the REX server using procedures outlined in the *System Administration Tasks* manual, and in the "NFS Configuration and Maintenance," "NIS Configuration and Maintenance," and "Remote Execution Facility (REX)" chapters of this manual?

3. Does the `inetd` security file (`/usr/adm/inetd.sec`) on the REX server limit access to the remote system for the `rexd` service?

4. Does the user have a user account on both the REX client and the REX server with matching UIDs?

5. Was `rexd` on the REX server started with the -r option? This causes access to be restricted based on `/etc/hosts.equiv` and the user's `.rhost` file on the REX server.

6. If the remote command is hung, is the NFS or NIS server down?

7. Is the problem associated with attempts to mount the file system containing the user's current working directory?

   a. Is the file system in the NFS server's `/etc/exports` file?

   b. Does the NFS server's `/etc/exports` entry for the file system restrict access to a specific netgroup or host?

Error Messages

The problem can exist on the server even though the error message may not occur on it.

Since most of the error messages are self-explanatory, you can determine the necessary corrective action when simple errors occur. For the other error messages, follow the corrective action supplied in the *HP-UX Reference* for that service. (These error messages are preceded by the name of the service.)

Errnos

NFS provides two `errno` values: `ESTALE` and `EREMOTE`.

`ESTALE` — You cannot reference the file because it no longer exists. This situation can occur since NFS allows a file opened by a client to be removed by a user on another node.

`EREMOTE` — You cannot mount file systems from a server that the server has remotely mounted (i.e., you cannot use NFS servers as NFS gateways).

Unsolved Problems

If you do not solve the problem after working through the previous troubleshooting steps and following flowcharts, call your HP support representative for assistance. Provide as much information about the problem as possible, including information from your network map and the following items:

- The activity you were attempting when the error occurred. Describe the HP-UX commands, job streams, result codes, and events leading to and including the problem.
- The version or update information for all software you are running. You should be able to find this information on your *install* or *update* media.
- The error messages you received. Record all error messages and numbers that appeared both on all nodes.
- The troubleshooting steps you tried.
- The problems you ruled out and why.

Flowchart Format

Each of the following flowcharts have a corresponding set of labelled explanations. You can use the flowcharts alone or with the explanatory text for more detail.

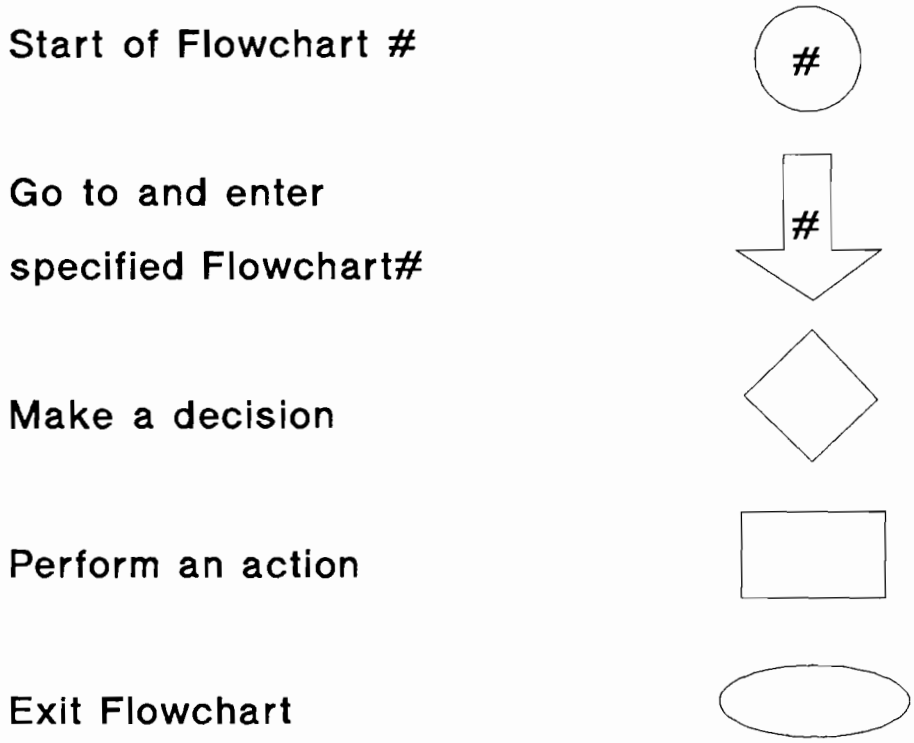

Flowchart Symbols

Note These flowcharts are for HP systems. Processes referenced in the flowcharts may not be part of NFS products from other vendors (e.g., `portmap`).

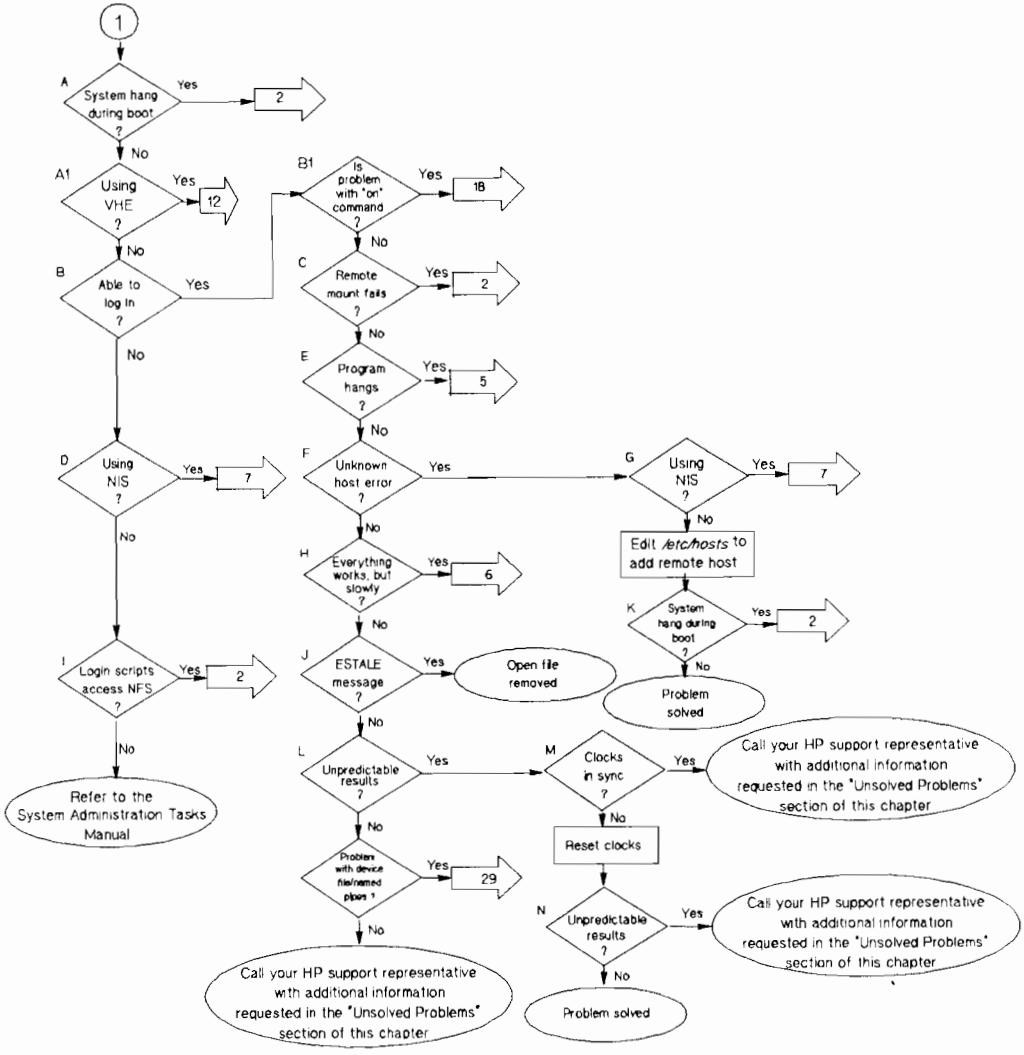

Flowchart 1: Initial Steps to Narrowing the Problem

Initial Steps to Narrowing the Problem (Flowchart 1)

Begin your troubleshooting efforts with Flowchart 1 since it helps you determine the best troubleshooting path based on the problem's symptoms.

Question	Yes: Action	No: Action
A. Does the system hang during boot when mounting remote files? Systems hanging during boot where remote mounts generally occur may indicate one or more servers are down or the network connection to one or more servers is faulty.	See Flowchart 2.	See A1.
A1. Are you using VHE?	See Flowchart 12.	See B.
B. Are you able to login?	See C.	You will receive error messages or the system will fail to respond if you cannot log into it. See D.
B1. Is the problem experienced while using the "on" command?	See Flowchart 18.	See C.
C. When trying to mount a remote file system, do error messages indicate the attempt failed?	See Flowchart 2.	See E.
D. Are you using NIS?	See Flowchart 7.	See I.
E. Do programs performing remote file accesses hang?	See Flowchart 5.	See F.

Question	Yes: Action	No: Action
F. Does the system report unknown host errors during execution of commands or programs?	See G.	See H.
G. Are you using NIS?	See Flowchart 7.	Edit `/etc/hosts` to add remote host, and then see K.
H. Does everything work, but slowly?	See Flowchart 6.	See J.
I. Do your login scripts perform NFS remote file accesses?	See Flowchart 2.	The problem is probably unassociated with the network services. Refer to the system login information in the *System Administration Tasks* manual.
M. Does the following message occur? `ESTALE`	The file was removed by another user. NFS allows file removal at any time.	See L.
K. Does the system hang during boot?	Restart Flowchart 1.	Problem solved.
L. Are you receiving unpredictable results when executing programs or commands?	See M.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
M. Are the server and client clocks synchronized?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Reset the clocks using the `date` command, and then see N.

Question	Yes: Action	No: Action
N. Do you receive unpredictable results to commands or programs?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Problem solved.

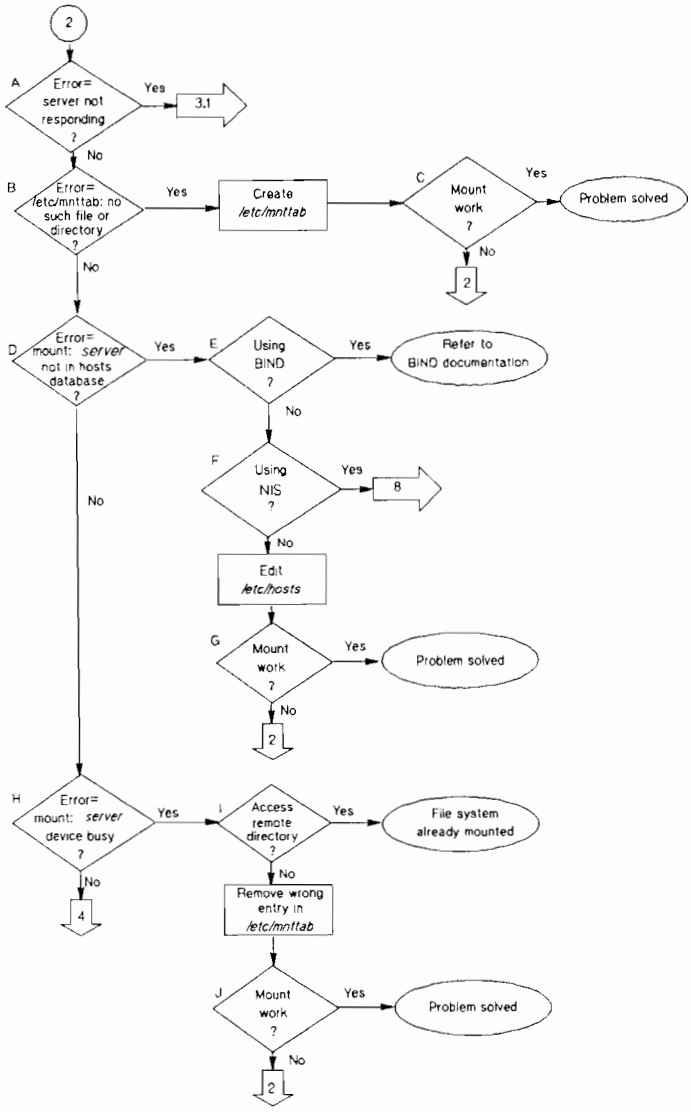

Flowchart 2: Mount Fails

Mount Fails (Flowchart 2)

Use Flowchart 2 if your system hangs during the booting process when remote file systems are mounted or if your remote mount attempts are unsuccessful.

Before using Flowchart 2, remember to check the `mount` command syntax and correct errors according to the error messages.

Question	Yes: Action	No: Action
A. Does the following error message occur on the client? `server not responding`	See Flowchart 3.1.	See B.
B. Does the following error message occur on the client? `/etc/mnttab: no such file or directory`	Create `/etc/mnttab` on the client, and then see C. The system uses `/etc/mnttab` to log all mounted file systems. *Note*: Generally, at boot time `/etc/rc` creates `/etc/mnttab`.	See D.
C. Can you mount the remote system?	Problem solved.	Restart Flowchart 2.
D. Does the following error message occur on the client? `mount: server not in hosts database`	See E.	See H.
E. Are you using BIND?	See the BIND documentation in *Installing and Administering ARPA Services*.	See F.

Question	Yes: Action	No: Action
F. Are you using NIS?	See Flowchart 8.	Edit `/etc/hosts` on the client to include the desired remote host, and then see G.
G. Can you mount the remote system?	Problem solved.	Restart Flowchart 2.
H. Does the following error message occur on the client? `mount: server device busy`	See I.	See Flowchart 4.
I. Can you access a remote directory in the desired remote file system?	You do not need to mount the file system since it is already mounted; problem solved.	On the client, remove the incorrect entry in `/etc/mnttab` for the remote file system you are trying to mount, and then see J.
J. Can you mount the remote system?	Problem solved.	Restart Flowchart 2.

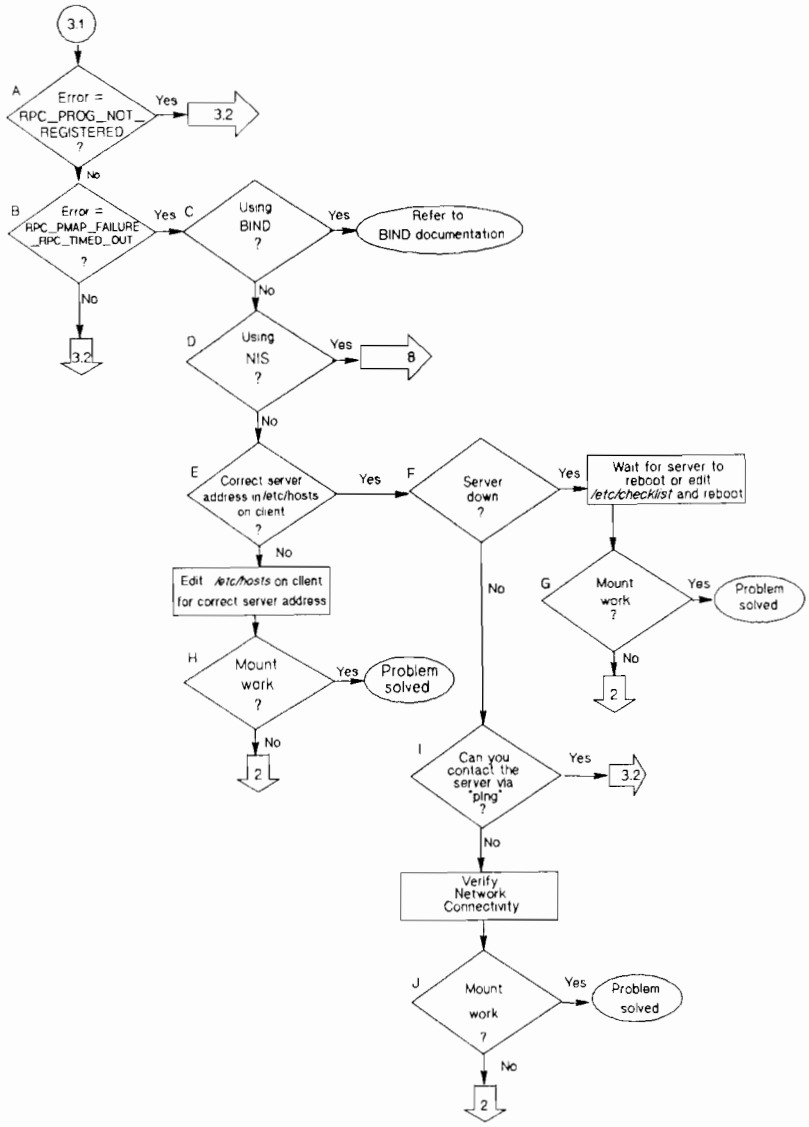

Flowchart 3.1: Server Not Responding

Server Not Responding (Flowchart 3.1)

This flowchart and corresponding instructions consist of two parts: Flowchart 3.1 and 3.2.

Question	Yes: Action	No: Action
A. Does the following error message occur? `RPC_PROG_NOT_REGISTERED`	See Flowchart 3.2.	See B.
B. Does the following error message occur? `RPC_PMAP_FAILURE: RPC_TIMED_OUT`	See C.	See Flowchart 3.2.
C. Are you using BIND?	See the BIND documentation in *Installing and Administering ARPA Services*.	See D.
D. Are you using NIS?	See Flowchart 8.	See E.
E. Is the server's address correct in the client's `/etc/hosts`?	See F.	Edit the client's `/etc/hosts` to include the correct address for the server you are trying to mount. See H.
F. Is the server you are trying to mount down? To check, ask your system administrator or try other network services to that system.	You have two options: - Do nothing on the system until the server reboots. - Edit the client's `/etc/checklist` to remove the NFS entry for that server; reboot the system. See G.	See I.

Question	Yes: Action	No: Action
G. Can you mount the remote system?	Problem solved.	See Flowchart 2.
H. Can you mount the remote system?	Problem solved.	See Flowchart 2.
I. Can you contact the server using the ping diagnostic? Refer to the *Installing and Administering LAN* manuals for `ping` diagnostic information	See Flowchart 3.2.	Refer to the *Installing and Administering LAN* manual to verify link connectivity, and then see J.
J. Can you mount the remote system?	Problem solved.	See Flowchart 2.

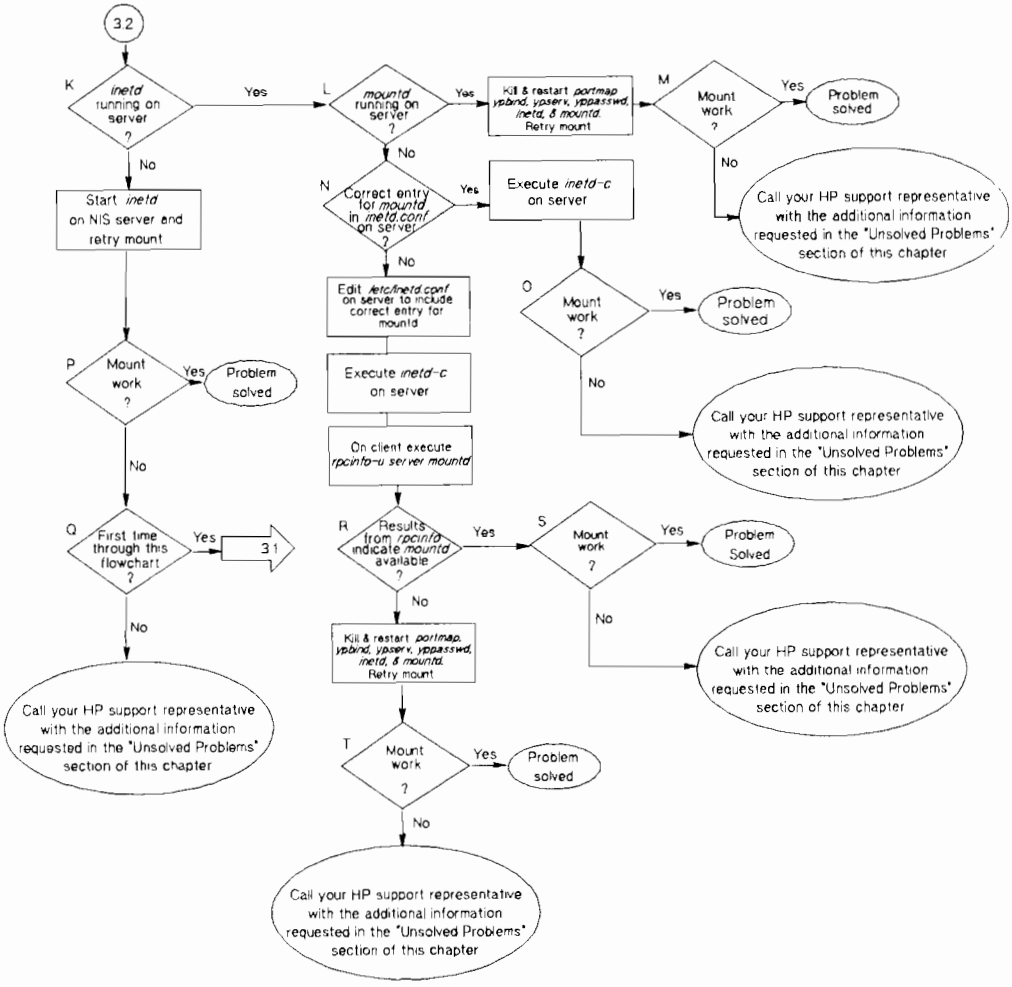

Flowchart 3.2: Server Not Responding

Server Not Responding (Flowchart 3.2)

Question	Yes: Action	No: Action
K. Is `inetd` running on the server?	See L.	Start `/etc/inetd` on the server, retry the mount, and then see P.
L. Is `mountd` running on the server?	Kill and restart the following daemons on the server in the order specified: - `portmap` - `ypbind` * - `ypserv` * - `yppasswdd` * - `inetd` - `mountd` * only if using NIS Retry the mount, and then see M.	See N.
M. Can you mount the remote system?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

Question	Yes: Action	No: Action
N. Is the correct `mountd` entry in `inetd.conf` on the server? Ensure the entry is not commented out with a # (pound sign).	Execute `inetd -c` on the server, and then see O.	1. Edit the server's `/etc/inetd.conf` file to include the correct `mountd` entry. 2. Execute `inetd -c` on the server to read changes in `/etc/inetd.conf`. 3. Execute `rpcinfo -u` on the client. `rpcinfo -u server mountd` 4. See R.
O. Can you mount the remote system?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
P. Can you mount the remote system?	Problem solved.	See Q.
Q. Is this the first time you used this flowchart for this problem?	Restart Flowchart 3.1.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

Question	Yes: Action	No: Action
R. Do the results from `rpcinfo -u` indicate a `mountd` process is available on the server?	See S.	Kill and restart the following daemons on the server in the order specified: - `portmap` - `ypbind` * - `ypserv` * - `yppasswdd` * - `inetd` - `mountd` * only if using NIS Retry the mount, and then see T.
S. Can you mount the remote system?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
T. Can you mount the remote system?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

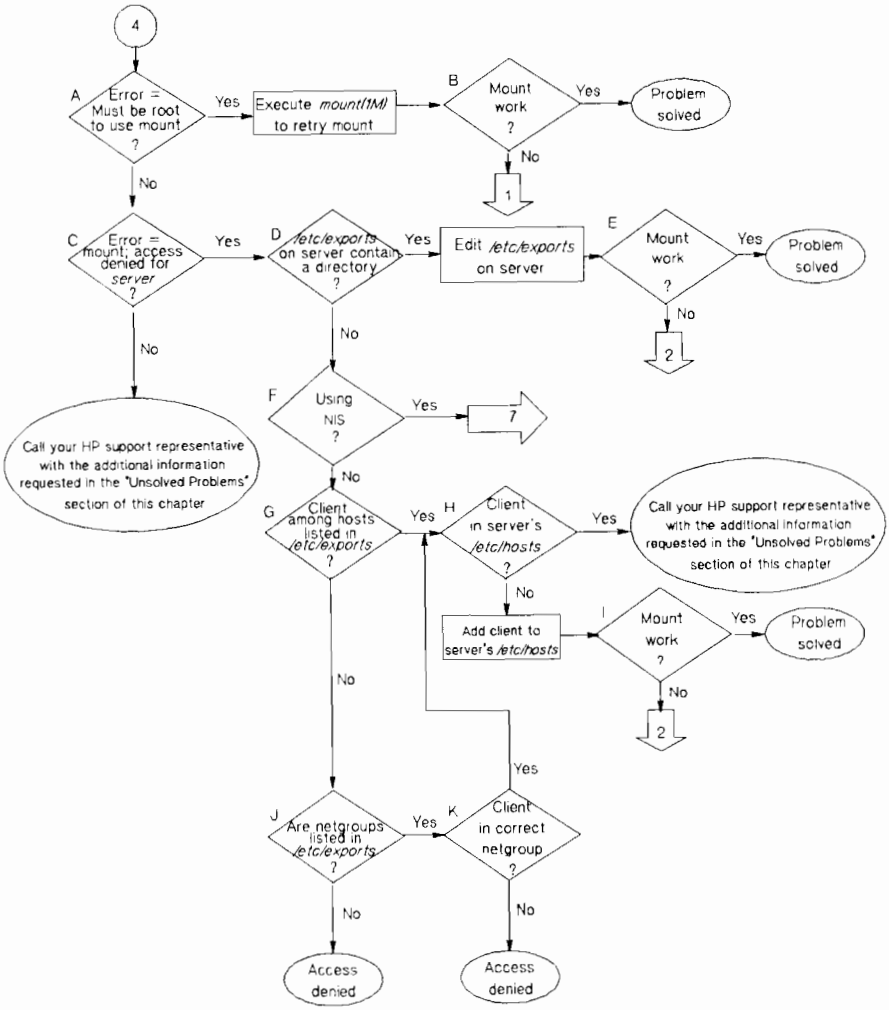

Flowchart 4: Restricted Access

Restricted Access (Flowchart 4)

Question	Yes: Action	No: Action
A. Does the following error message occur on the client? `Must be root to use mount`	Log in as superuser, execute `mount`, and then see B.	See C.
B. Can you mount the remote system?	Problem solved.	See Flowchart 1.
C. Does the following error message occur on the client? `mount: access denied for server`	See D.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
D. Does the server's `/etc/exports` file list a directory rather than a file system?	Edit the server's `/etc/exports` to contain the file system rather than a directory, and then see E.	See F.
E. Can you mount the remote system?	Problem solved.	See Flowchart 2.
F. Are you using NIS?	See Flowchart 7.	See G.
G. If hosts are listed in `/etc/exports`, is the client among the hosts listed for the desired file system?	See H.	See J.
H. Is the client listed in the server's `/etc/hosts`?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Add client to server's `/etc/hosts`, and then see I.
I. Can you mount the remote system?	Problem solved.	See Flowchart 2.

Question	Yes: Action	No: Action
J. Are netgroups listed for this file system in server's `/etc/exports`?	See K.	Access for this client is deliberately denied.
K. Is the client listed in the appropriate netgroup for this file system in `/etc/netgroup`?	See H.	Access for this client is deliberately denied.

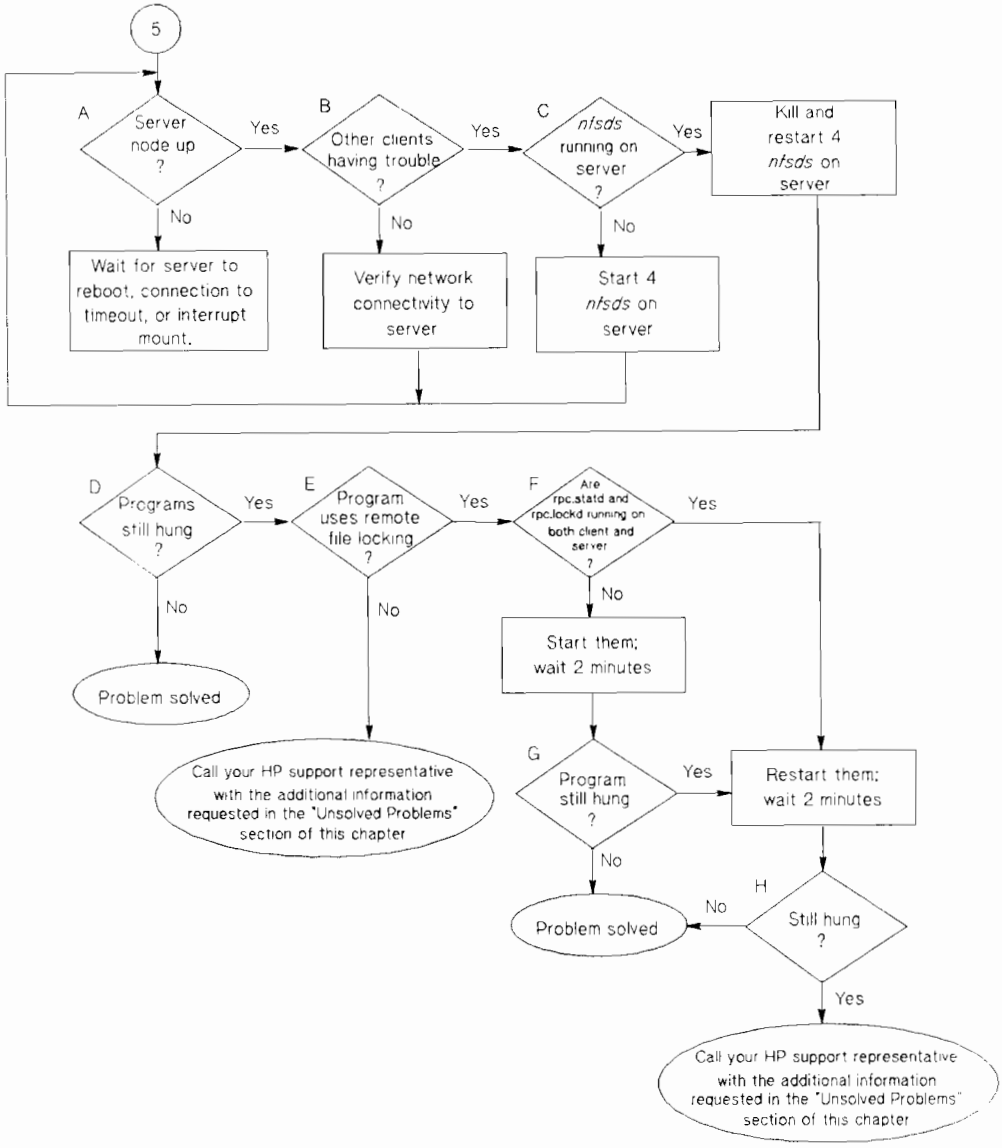

Flowchart 5: Programs Hang

Programs Hang (Flowchart 5)

Programs are most likely to hang if network communication is lost to the server, if the server is down, or if daemons are hung.

Question	Yes: Action	No: Action
A. Is the server node running?	See B.	For hard mounts, do *one* of the following: - Wait for the server to reboot. - Interrupt the mount. For soft mounts, wait for the mount to time out. See A.
B. Are other client nodes having trouble?	See C.	Verify the network connectivity. Refer to the *Installing and Administering LAN* manuals. See A.
C. Are `nfsd` daemons running on the server?	Kill and restart four `nfsd` daemons on the server, and then see D.	Start four `nfsd` daemons on the server, and then see A.
D. Do the programs hang?	See E.	Problem solved.
E. Does the program use remote file locking?	See F.	Call your HP Support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

Question	Yes: Action	No: Action
F. Are `rpc.statd` and `rpc.lockd` running on both the client and the server?	Restart them and wait 2 minutes. See H.	Start them and wait 2 minutes. See G.
G. Does the program still hang?	Restart rpc.statd and rpc.lockd on both the client and the server. Wait 2 minutes. See H.	Problem solved.
H. Does the program still hang?	Call your HP Support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Problem solved.

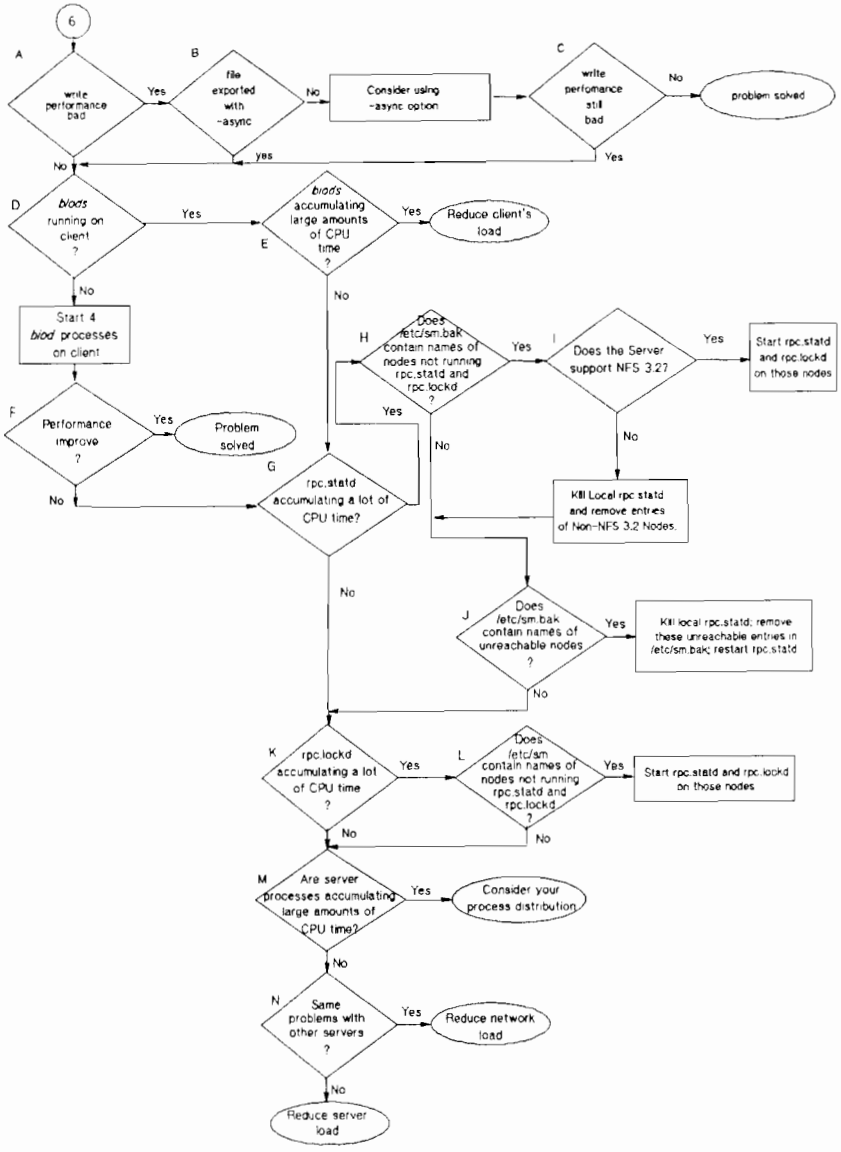

Flowchart 6: Performance Problems

Performance Problems (Flowchart 6)

Question	Yes: Action	No: Action
A. Is the write performance bad?	See B.	See D.
B. Is the file exported with `-async`?	See D	Consider using `-async` option, then see C.
C. Is the write performance still bad?	See D.	Problem Solved.
D. Are the biods running on the client?	See E.	Start four `biod` processes on the client, and then see F.
E. Are the client `biod` daemons accumulating large amounts of CPU time? 1. List the client processes using `ps`. 2. Copy a large file to the server system, and list the client `biod` processes again. 3. Compare the CPU time for the `biod` processes before and after the file copy.	Reduce the client's load to fewer NFS transactions by reducing the number of users or storing more files locally.	See G.
F. Has the performance improved?	Problem solved.	See G.
G. Is `rpc.statd` accumulating a lot of CPU time? (On the client?)	See H.	See K.
H. Does `/etc/sm.bak` contain names of nodes not running `rpc.statd` and `rpc.lockd`?	See I.	See J.

Question	Yes: Action	No: Action
I. Does the server support NFS 3.2 functionality? (HP-UX 6.5 or later for the Series 300/400. HP-UX 7.0 or later for other HP architectures.)	Start `rpc.statd` and `rpc.lockd` on those nodes.	Kill local `rpc.statd` and remove entries of non-NFS 3.2 nodes, then see J.
J. Does `/etc/sm.bak` contain names of unreachable nodes?	Kill local `rpc.statd`, remove these unreachable entries in `/etc/sm.bak`, and restart `rpc.statd`	See K.
K. Is `rpc.lockd` accumulating a lot of CPU time? (On the client?)	See L.	See M.
L. Does `/etc/sm.bak` contain names of nodes not running `rpc.statd` and `rpc.lockd`?	Start `rpc.statd` and `rpc.lockd` on those nodes.	See M.
M. Are processes on the server accumulating large amounts of CPU time (especially `nfsd`, `inetd`, and `portmap`)?	Consider whether you need to distribute your processing by adding additional systems.	See N.
N. Are the same performance problems evident with other servers?	Reduce the network load.	Reduce the server's load by adding more servers.

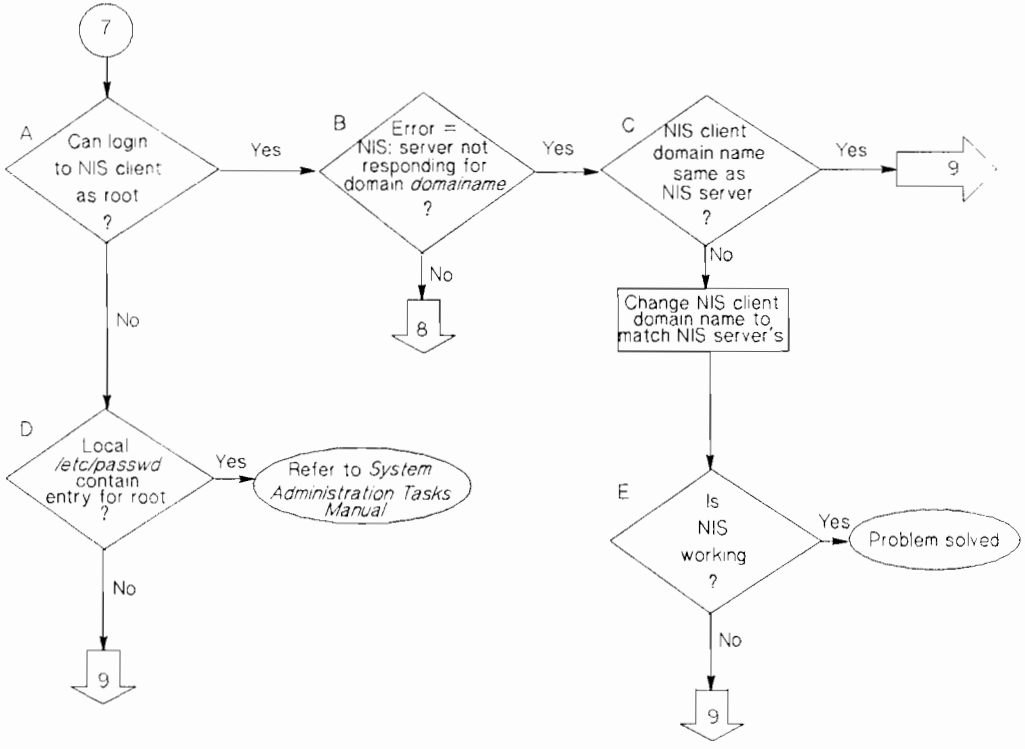

Flowchart 7: Initial Steps to Troubleshooting NIS

Initial Steps to Troubleshooting NIS (Flowchart 7)

Question	Yes: Action	No: Action
A. Can you login as root on the NIS client?	See B.	See D.
B. Does the following error message occur on the console or in the `ypbind` log file? `NIS: server not responding for domain domain_name`	See C.	See Flowchart 8.
C. Is the NIS client's NIS domain name the same as the NIS server's?	See Flowchart 9.	Change the NIS client's NIS domain name to be the same as the NIS server's, and then see E. `domainname domain_name`
D. Does the local `/etc/passwd` file contain an entry for root?	The problem is not associated with NIS or NFS. Refer to the *System Administration Tasks Manual*.	You cannot log into the NIS client until NIS is functioning unless you have an entry for a user in the local `/etc/passwd` file. See Flowchart 9.
E. Is NIS working? If you can access the NIS server's maps using `ypcat` or `ypmatch`	NIS is probably functioning correctly.	Problem solved.

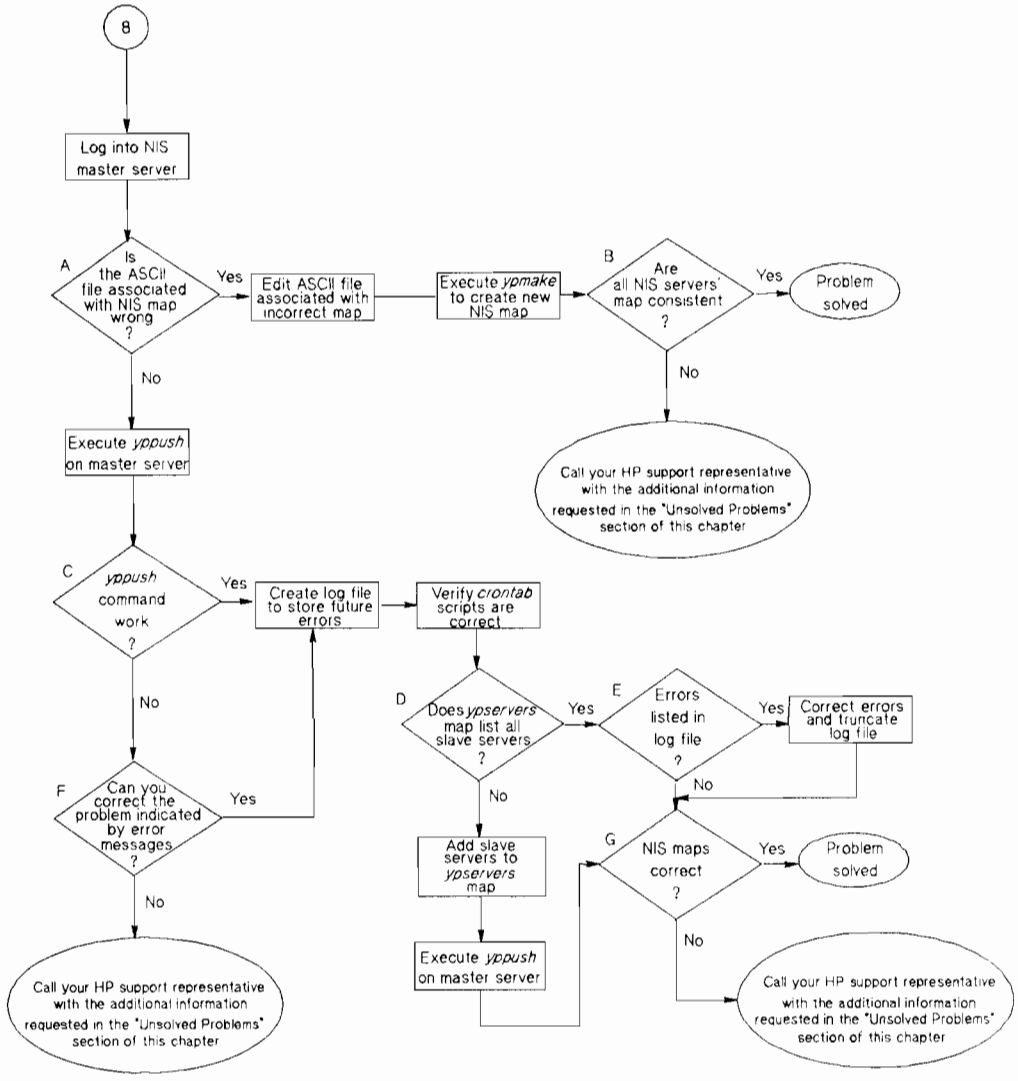

Flowchart 8: Incorrect NIS Maps

Incorrect NIS Maps (Flowchart 8)

Log into the NIS master server as root before starting Flowchart 8.

Question	Yes: Action	No: Action
A. On the NIS master server, does the ASCII file associated with the NIS map need to be updated (e.g., update `/etc/hosts`)?	1. Edit the ASCII file associated with the incorrect NIS map. 2. Execute `ypmake` to create and distribute a new map to the NIS slave servers. 3. See B.	Execute `yppush` on the NIS master server, and then see C. `yppush map_name`
B. Are all NIS server's maps consistent? You can determine this by executing `yppoll` and then comparing order numbers.	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
C. Does `yppush` work correctly? If you do not receive error messages associated with the command, it probably executed successfully.	1. Create the log file `/usr/etc/yp/ypxfr.log` to trap future errors associated with `yppush` on each NIS slave server. 2. Verify that `crontab` scripts (on each slave server) copying the maps are correct. 3. See D.	See F.

Question	Yes: Action	No: Action
D. Does the `ypservers` map list all NIS slave servers? `ypcat -k ypservers`	See E.	1. Add any missing NIS slave server to the `ypservers` map. 2. Execute `yppush` on the NIS master server to update all NIS slave servers. 3. See G.
E. Does `/usr/etc/yp/ypxfr.log` on the slave server list errors?	Correct the errors, truncate the log file, and then see G.	See G.
F. Can you correct the problem indicated by the error message?	1. Create the log file `/usr/etc/yp/ypxfr.log` to trap future errors associated with `yppush` on each NIS slave server. 2. Verify that `crontab` scripts (on each slave server) distributing the maps are correct. 3. See D.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
G. Are the NIS maps correct?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

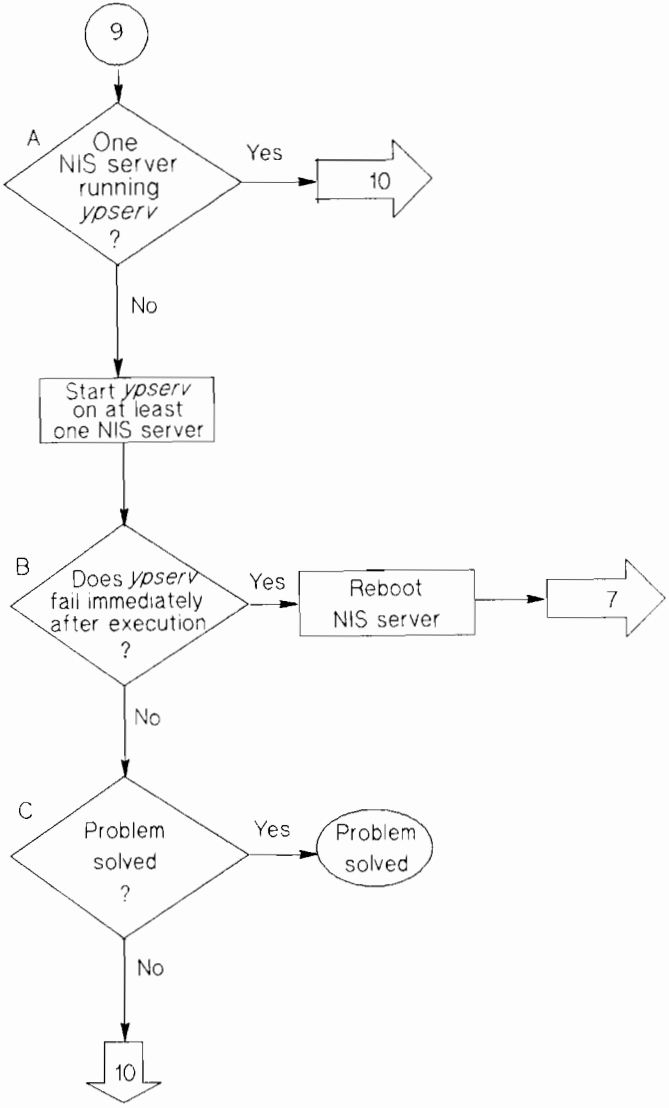

Flowchart 9: ypserv Problems

ypserv Problems (Flowchart 9)

Question	Yes: Action	No: Action
A. Is at least one NIS server in the NIS domain running `ypserv`?	See Flowchart 10.	Start `ypserv` on at least one NIS server in the NIS domain, and then see B.
B. Does `ypserv` fail immediately after starting it?	Reboot the NIS server, and then see Flowchart 7.	See C.
C. Is the problem solved?	Problem solved.	See Flowchart 10.

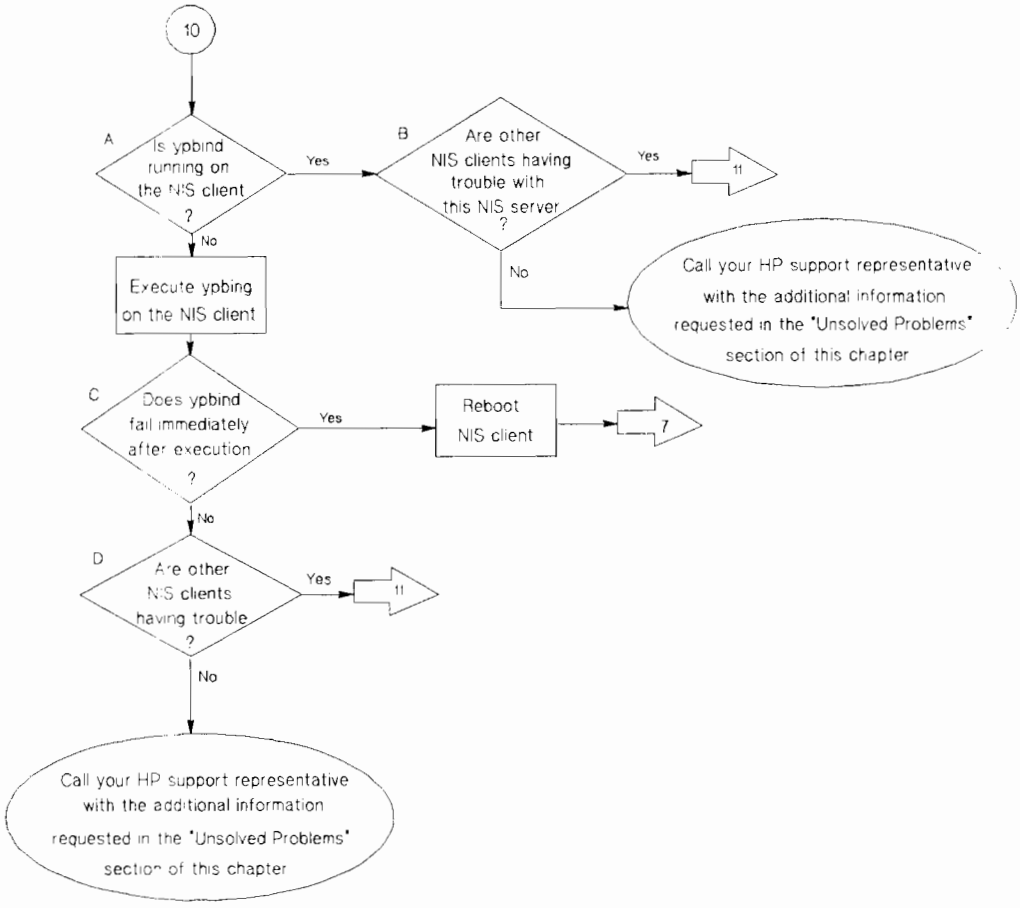

Flowchart 10: ypbind Problems

ypbind Problems (Flowchart 10)

Question	Yes: Action	No: Action
A. Is `ypbind` running on the NIS client?	See B.	Execute `ypbind` on the NIS client, and then see C.
B. Are other NIS clients having trouble with this NIS server?	See Flowchart 11.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
C. Does `ypbind` crash immediately after starting it?	Reboot the NIS client, and then see Flowchart 7.	See D.
D. Are other NIS clients having trouble with this NIS server?	See Flowchart 11.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

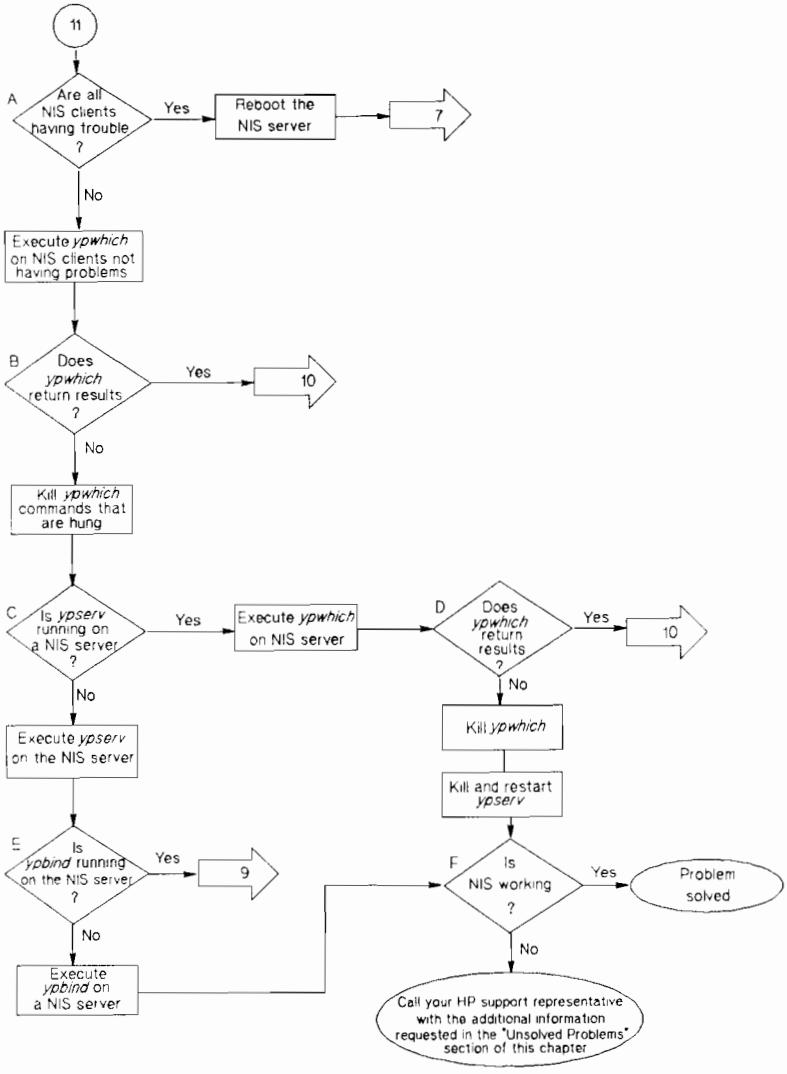

Flowchart 11: Multiple NIS Client Problems

Multiple NIS Client Problems (Flowchart 11)

Question	Yes: Action	No: Action
A. Are all NIS clients having trouble with this NIS server?	Reboot the NIS server, and then see Flowchart 7.	Execute `ypwhich` on the NIS client nodes not having problems, and then see B.
B. Does the `ypwhich` command return results on the NIS client?	See Flowchart 10.	Kill `ypwhich` commands that are hung on NIS clients, and then see C.
C. Is `ypserv` running on the NIS server?	Execute `ypwhich` on the NIS server, and then see D.	Execute `ypserv` on the NIS server, and then see E.
D. Does `ypwhich` return results on the NIS server?	See Flowchart 10.	1. Kill `ypwhich` on the NIS server. 2. Kill and restart `ypserv`. 3. See F.
E. Is `ypbind` running on the NIS server?	See Flowchart 9.	Execute `ypbind` on the NIS server, and then see F.
F. Is NIS functioning correctly on all NIS clients?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

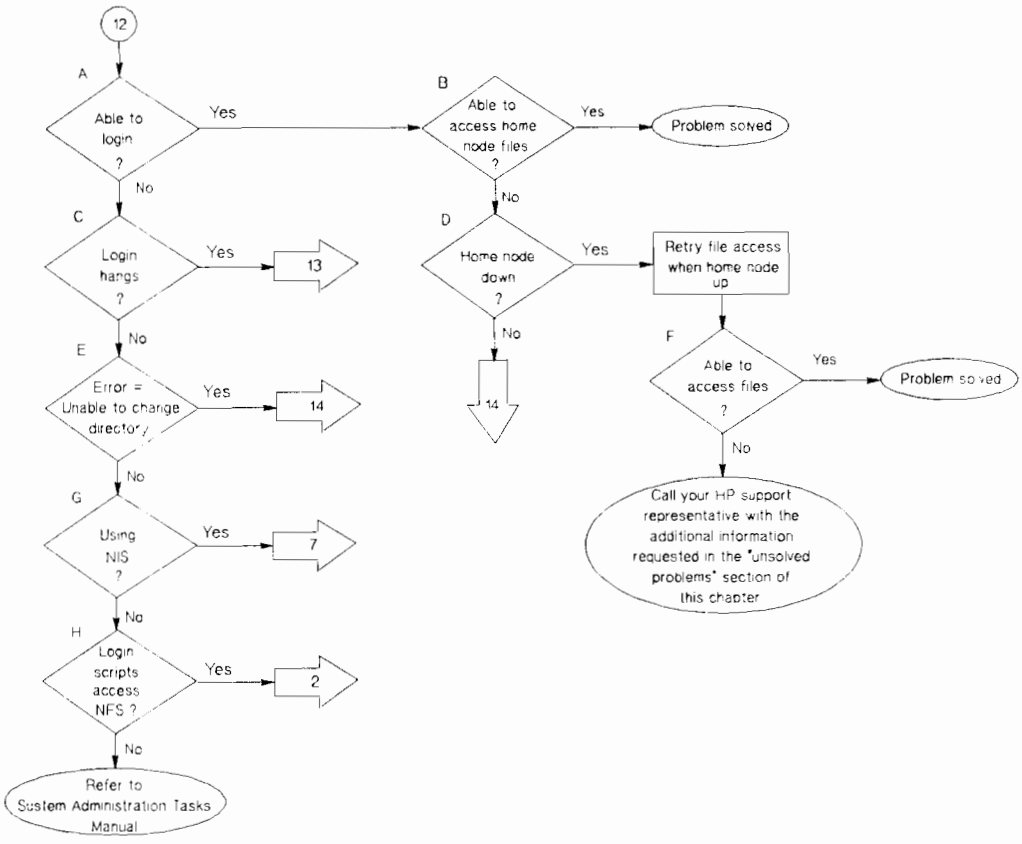

Flowchart 12: Initial Steps to Troubleshooting VHE

Initial Steps to Troubleshooting VHE (Flowchart 12)

Question	Yes: Action	No: Action
A. Are you able to log in?	See B.	See C.
B. Are you able to access files on the home node?	No problem.	See D.
C. Does the machine hang during login?	See Flowchart 13.	See E.
D. Is the home node down?	Retry accessing files when the home node is up; then see F.	See Flowchart 14.
E. Do you receive the following error message? `Unable to change directory to home directory`	See Flowchart 14.	See G.
F. Are you able to access files on the home node?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
G. Are you using NIS?	See Flowchart 7.	See H.
H. Do your login scripts perform NFS remote file access?	See Flowchart 2.	The problem is probably unassociated with the network services. Refer to the system login information in the *System Administration Tasks* manual.

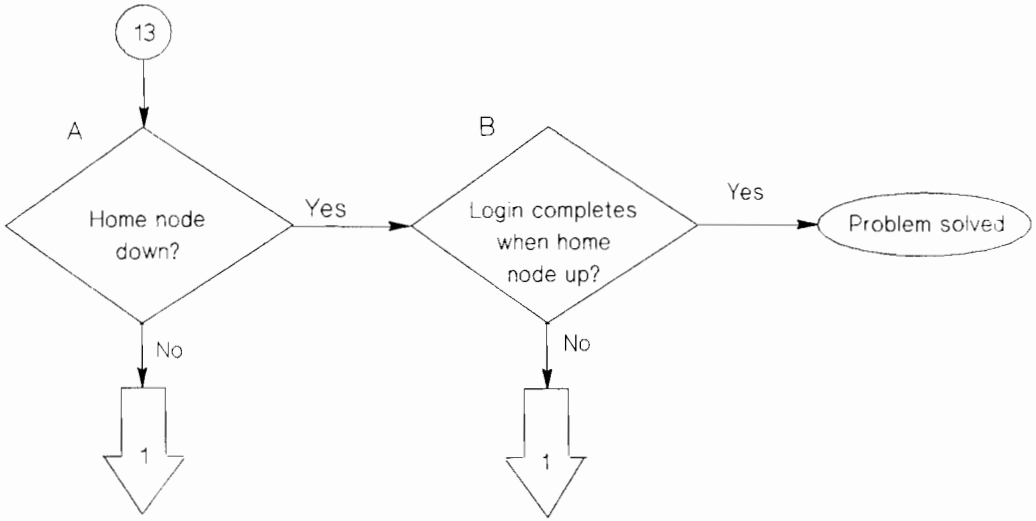

Flowchart 13: Home Node Goes Down After Mount Complete

Home Node Goes Down After Mount Complete (Flowchart 13)

Question	Yes: Action	No: Action
A. Is the home node down?	Try logging in again once the home node comes up; then see B.	See Flowchart 1.
B. Does the login complete once the home node comes up?	Problem solved.	See Flowchart 1.

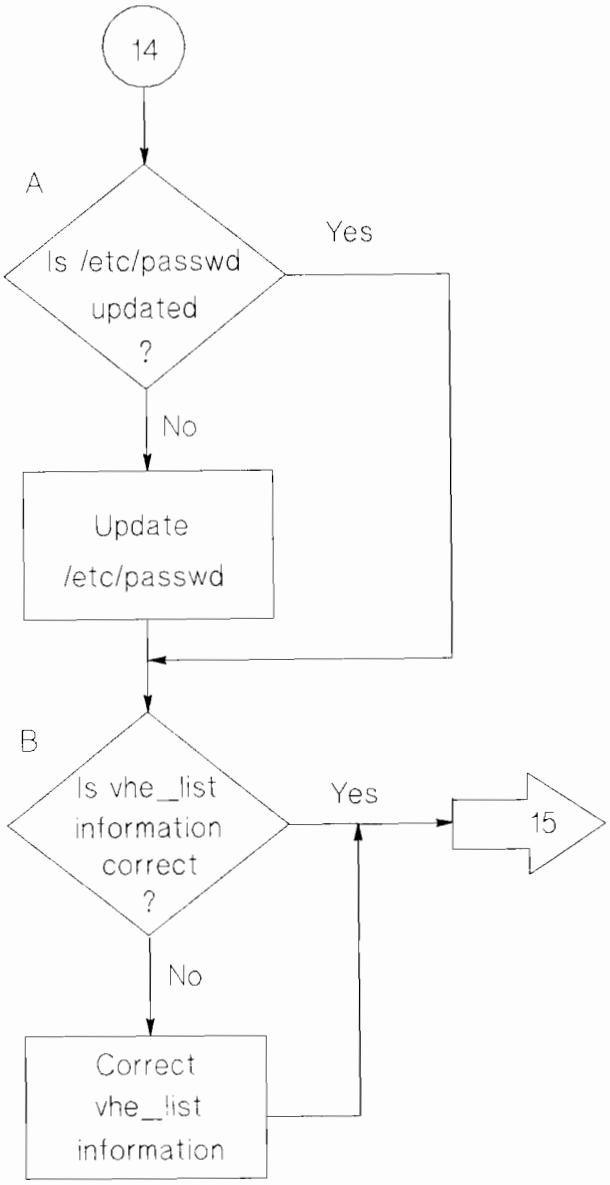

Flowchart 14: Checking /etc/passwd and /etc/vhe_list Files

Checking /etc/passwd and /etc/vhe_list Files (Flowchart 14)

Question	Yes: Action	No: Action
A. Is the `/etc/passwd` file updated to prefix the home directory with the NFS mount point?	See B.	Update the `/etc/passwd` file as described in the "VHE Configuration and Maintenance" chapter; go to B.
B. Is the information in the `/etc/vhe_list` file correct?	See Flowchart 15.	Correct the `/etc/vhe_list` file information; see Flowchart 15.

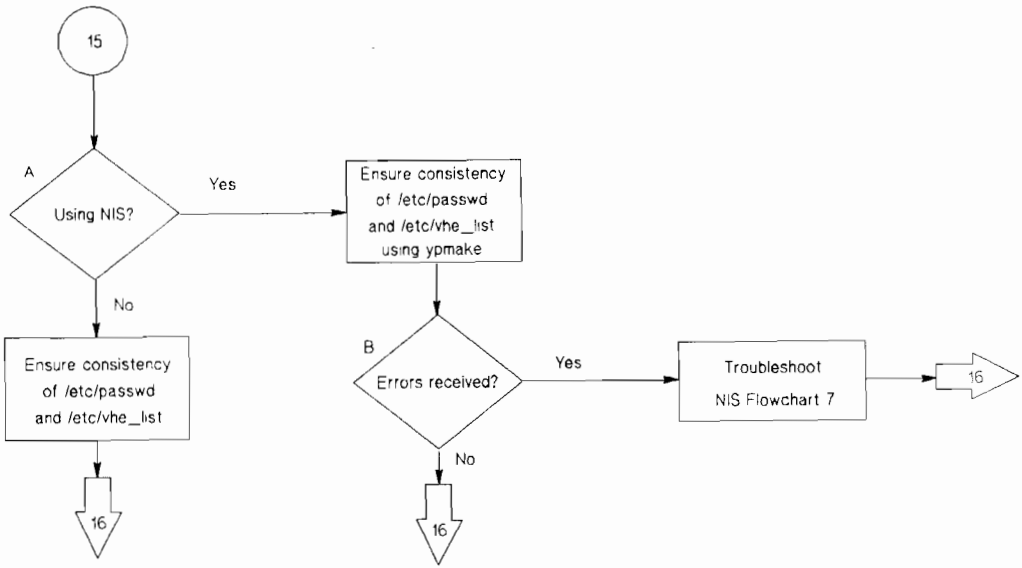

Flowchart 15: Consistency of /etc/passwd and /etc/vhe_list

Consistency of /etc/passwd and /etc/vhe_list (Flowchart 15)

Question	Yes: Action	No: Action
A. Are you using the Network Information Service (NIS) to ensure consistency of `/etc/passwd` and `/etc/vhe_list` information?	Ensure consistency of the `/etc/passwd` and `/etc/vhe_list` files on all nodes in the VHE group by executing the following command: `/usr/etc/yp/ypmake passwd vhe_list` See B.	Ensure consistency of the `/etc/passwd` and `/etc/vhe_list` files on all nodes in the VHE group. See Flowchart 16.
B. Did you receive any errors when executing `ypmake`?	Go to the NIS Flowchart 7 and complete troubleshooting steps; then return to VHE Flowchart 16.	Go to Flowchart 16.

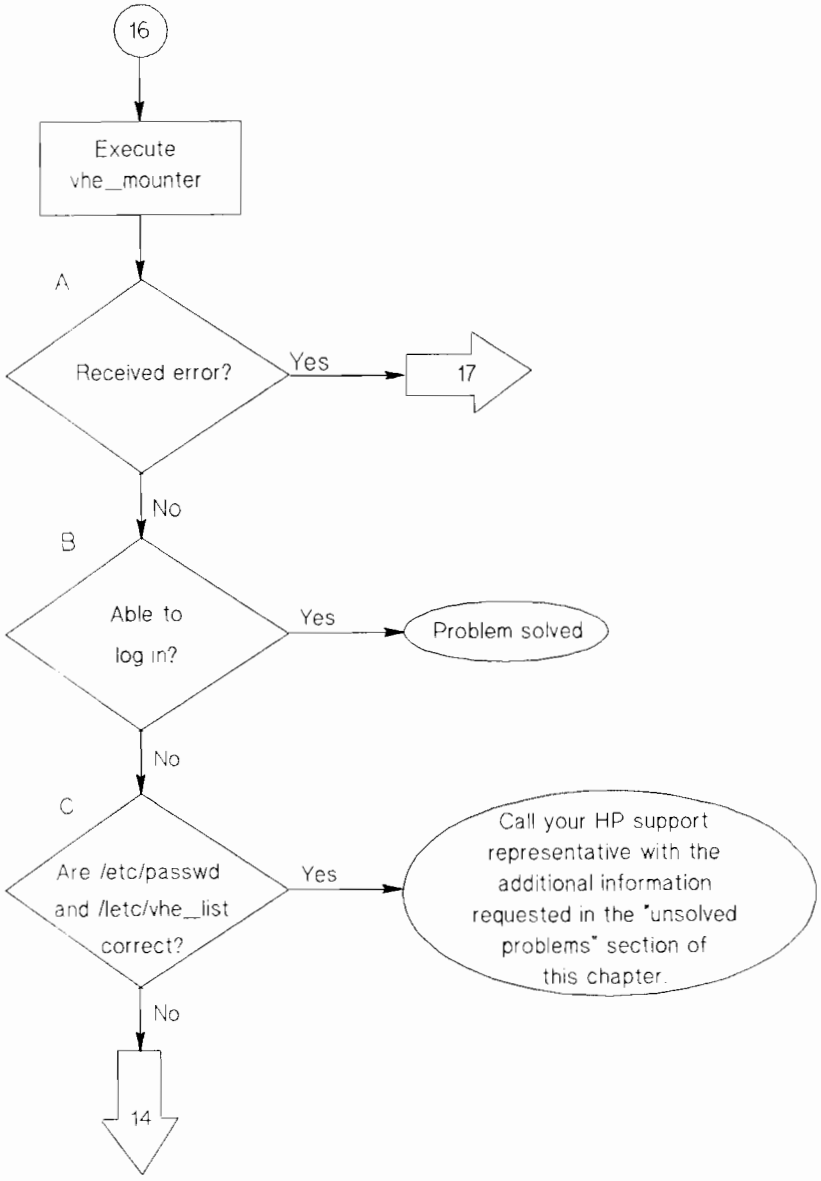

Flowchart 16: Execution of vhe-mounter

Execution of vhe_mounter (Flowchart 16)

Question	Yes: Action	No: Action
A. Did you receive any errors while executing `vhe_mounter`?	See Flowchart 17.	See B.
B. Are you able to log in?	Problem Solved.	See C.
C. Is the information for the home node entered into the `/etc/passwd` and `/etc/vhe_list` files?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 14.

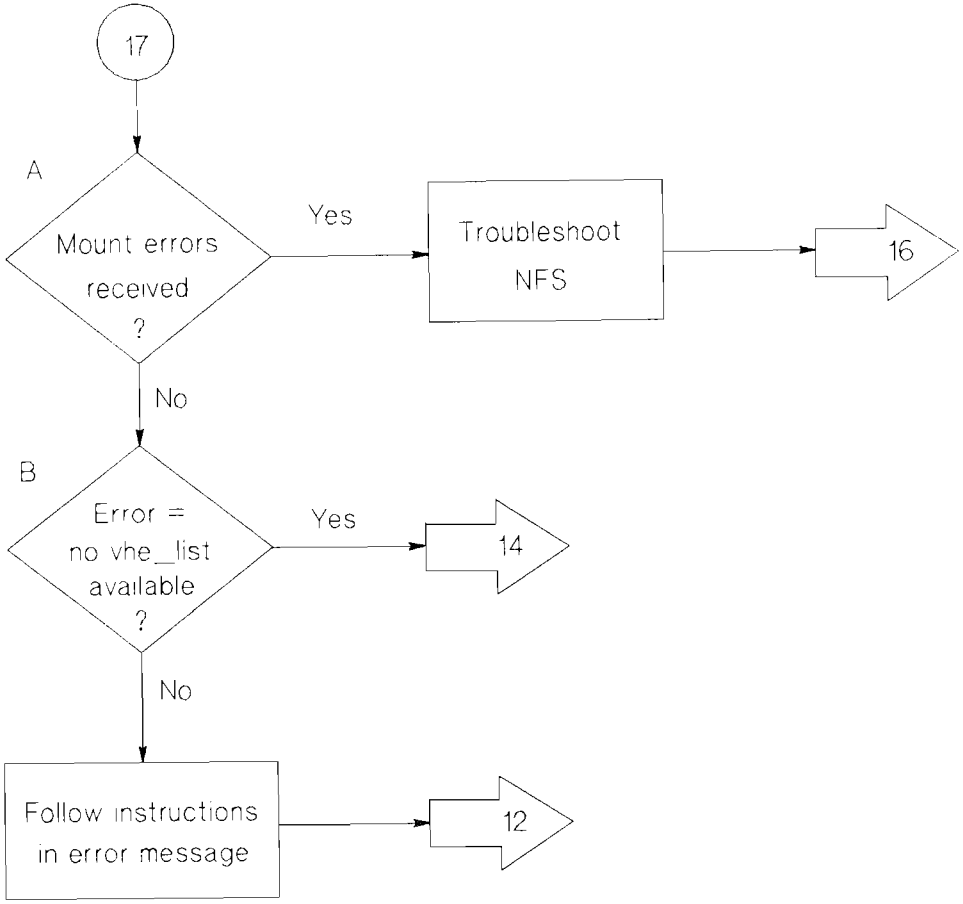

Flowchart 17: Error Message from vhe_mounter

Error Message from vhe_mounter (Flowchart 17)

Question	Yes: Action	No: Action
A. Were any mount errors encountered (mount errors begin with `mount:`)?	Troubleshoot NFS (Flowchart 1); then see Flowchart 16.	See B.
B. Does the following error message occur? `no vhe_list available`	See Flowchart 14.	If an error message other than those mentioned is printed, follow the instructions in that error message; then re-enter Flowchart 12 to see if problem is solved.

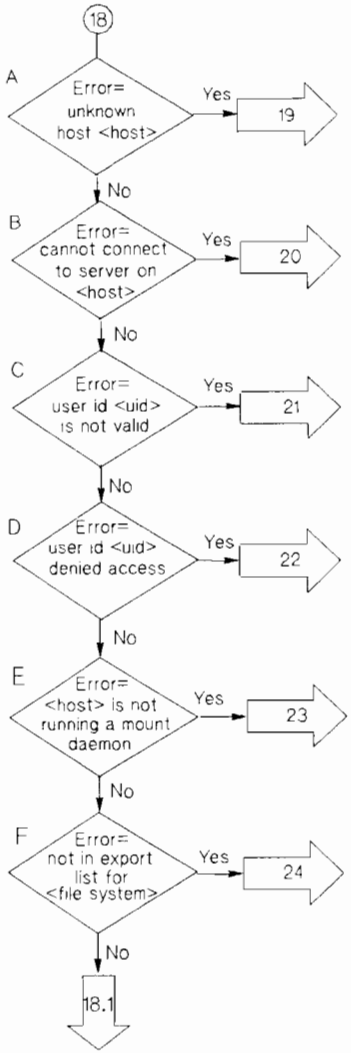

Flowchart 18: Initial Steps to Troubleshoot REX

Initial Steps to Troubleshoot REX (Flowchart 18)

Question	Yes: Action	No: Action
A. Does the following error message appear? on: unknown host <host>	See Flowchart 19.	See B.
B. Does the following error message appear? on: cannot connect to server on <host>	See Flowchart 20.	See C.
C. Does the following error message appear? on: rexd: user id is not valid	See Flowchart 21.	See D.
D. Does the following error message appear? on <server>: rexd: user id <uid> denied access	See Flowchart 22.	See E.
E. Does the following error message appear? on: <server>rexd: <host> is not running a mount daemon	See Flowchart 23.	See F.
F. Does the following error message appear? on <server>: rexd: not in export list for filesystem	See Flowchart 24.	See Flowchart 18.1.

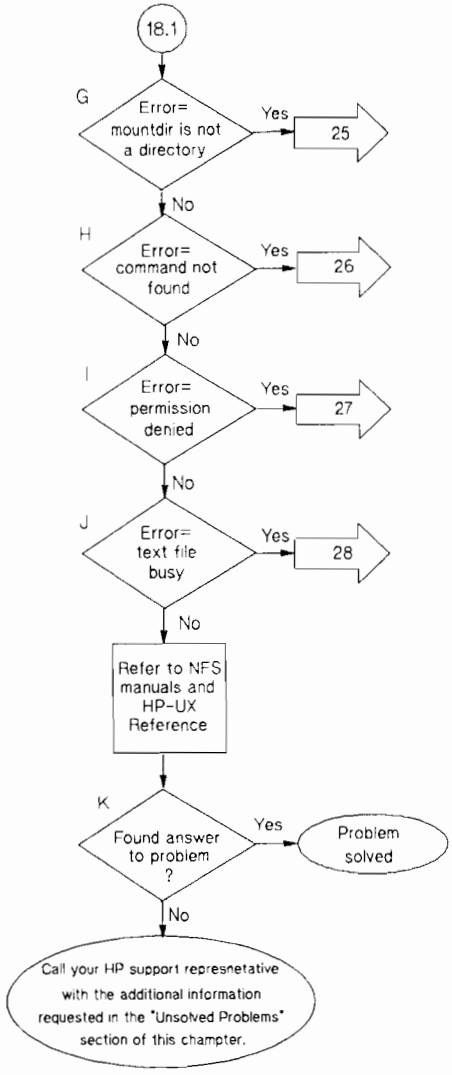

Flowchart 18.1: Initial Steps to Troubleshoot REX

Initial Steps to Troubleshoot REX (Flowchart 18.1)

Question	Yes: Action	No: Action
G. Does the following error message appear? `on <server>: rexd: (<mountdir>) is not a directory`	See Flowchart 25.	See H.
H. Does the following error message appear? `on <server>: rexd: command not found`	See Flowchart 26.	See I.
I. Does the following error message appear? `on <server>: rexd: permission denied`	See Flowchart 27.	See J.
J. Does the following error message appear? `on <server>: rexd: text file busy`	See Flowchart 28.	Refer to NFS manuals and *HP-UX Reference*. See K.
K. Found answer to your problem?	Problem solved.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

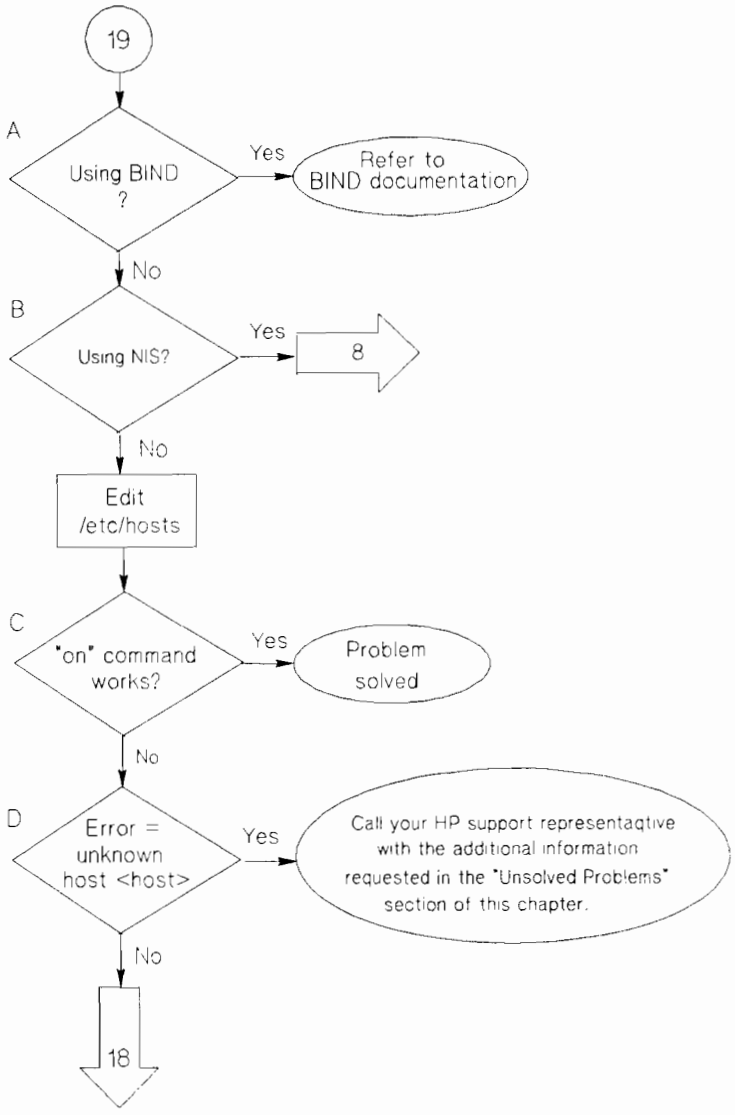

Flowchart 19: Unknown Host

Unknown Host (Flowchart 19)

Question	Yes: Action	No: Action
A. Is your node using BIND?	Refer to BIND documentation in *Installing and Administering ARPA Services*.	See B.
B. Is your node using the Network Information Service (NIS)?	See Flowchart 8.	Edit `/etc/hosts` on the client to include the desired remote host. See C.
C. `on` command works now?	Problem solved.	See D.
D. Does the following error message appear? `on: unknown <host>`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

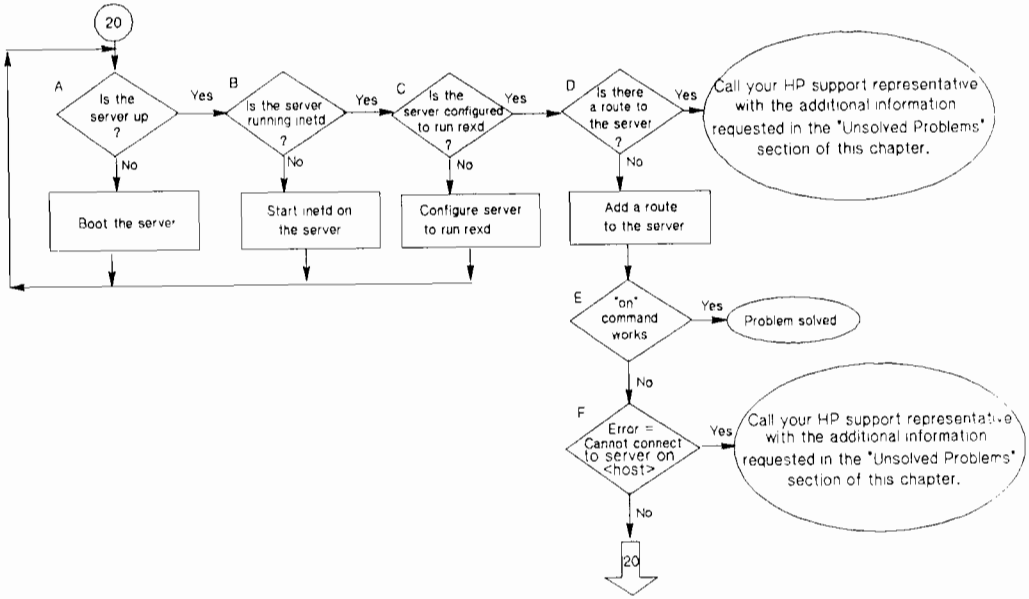

Flowchart 20: Cannot Connect to REX Server

Cannot Connect to REX Server (Flowchart 20)

Question	Yes: Action	No: Action
A. Is the rex server node up?	See B.	Boot the rex server node. See A.
B. Is the rex server node running `inetd`?	See C.	Start `inetd` on the rex server node. See A.
C. Is the rex server configured to run `rexd`?	See D.	Configure the rex server to run `rexd` by editing `/etc/inetd.conf` on the rex server, uncommenting the `rpc.rexd` line, and issuing the `inetd -c` command. See A.
D. Is there a route to the rex server?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Add a route using the `route` command. See E.
E. `on` command works now?	Problem solved.	See F.
F. Does the following error message appear? `on: cannot connect to server on <host>`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

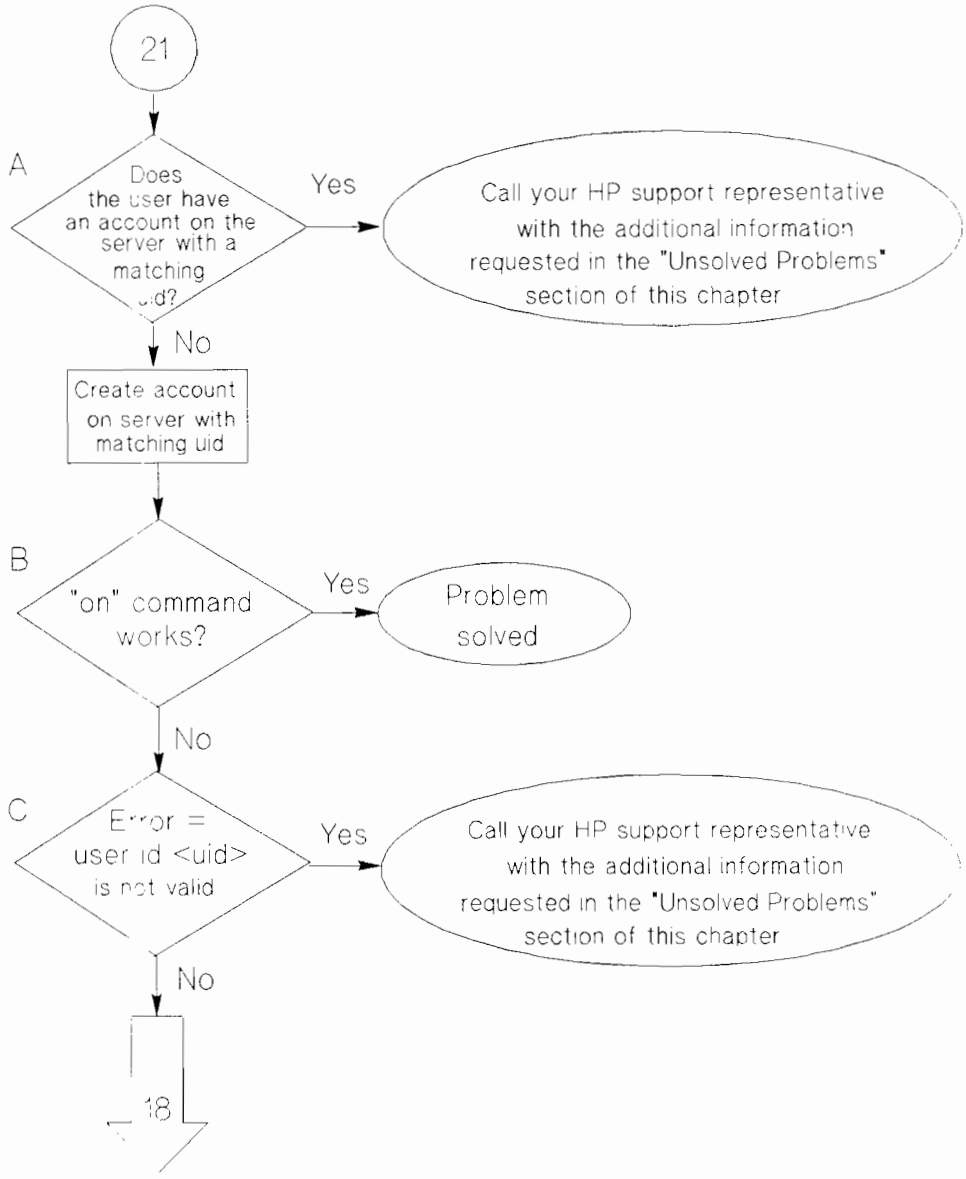

Flowchart 21: User ID Not Valid

User ID Not Valid (Flowchart 21)

Question	Yes: Action	No: Action
A. Does the user have an account on the rex server with a uid which matches the user's uid on the client?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Create an account on the rex server for the user with a matching uid. See B.
B. `on` command works now?	Problem solved.	See C.
C. Does the following error message appear? `on: rexd: user id is not valid`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

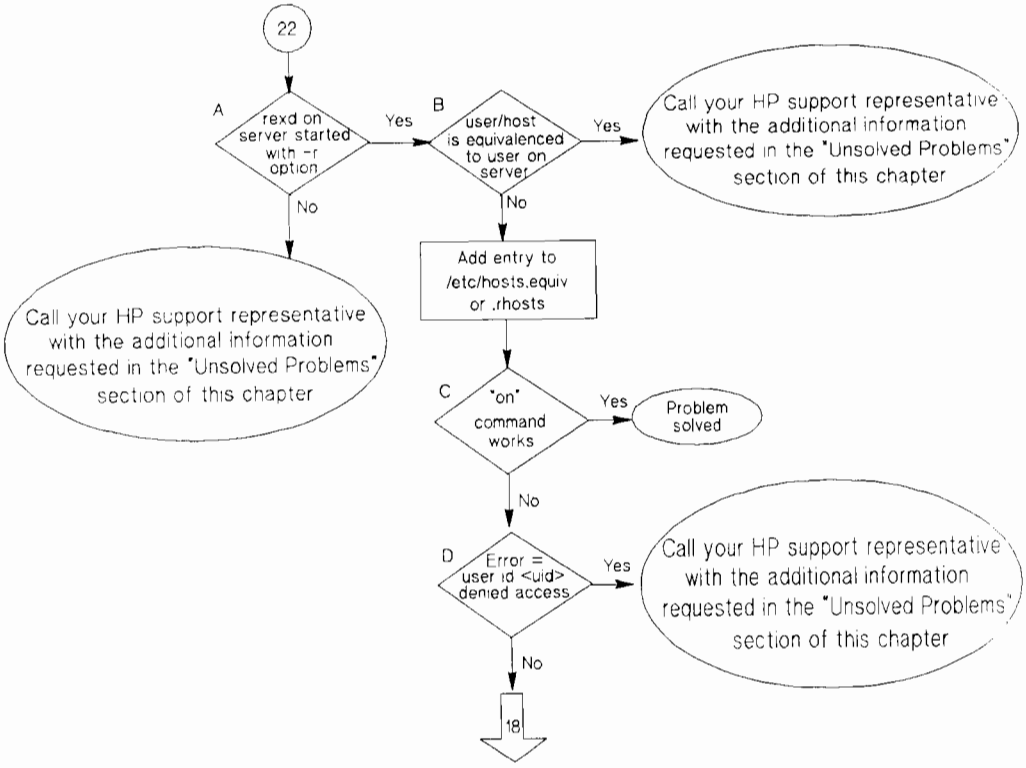

Flowchart 22: User ID Denied Access

User ID Denied Access (Flowchart 22)

Question	Yes: Action	No: Action
A. `Rexd` or rex server started with `-r` option?	See B.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
B. User or client host is equivalenced by entry in `.rhosts` or `/etc/hosts.equiv` file?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Edit `.rhosts` or `/etc/hosts.equiv` file to add an entry for the user or the client host. See C.
C. `on` command works now?	Problem solved.	See D.
D. Does the following error message appear? `on <server>: rexd: user id <uid> denied access`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

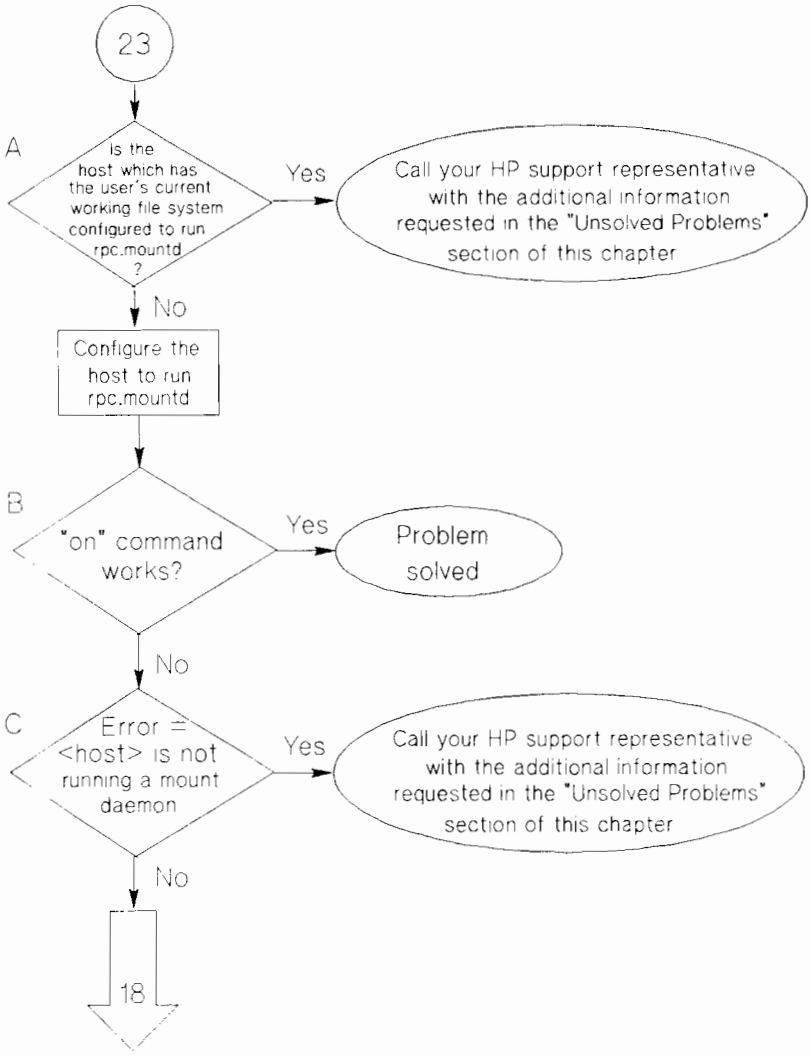

Flowchart 23: REX Server Not Running Mount Daemon

REX Server Not Running Mount Daemon (Flowchart 23)

Question	Yes: Action	No: Action
A. Is the host which has the user's current working file system physically mounted configured to run `rpc.mountd`?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Configure the NFS server node to run `rpc.mountd` by editing its `/etc/inetd.conf` and executing `inetd -c`. See B.
B. `on` command works now?	Problem solved.	See C.
C. Does the following error message appear? `on: <server>rexd: <host> is not running a mount daemon`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

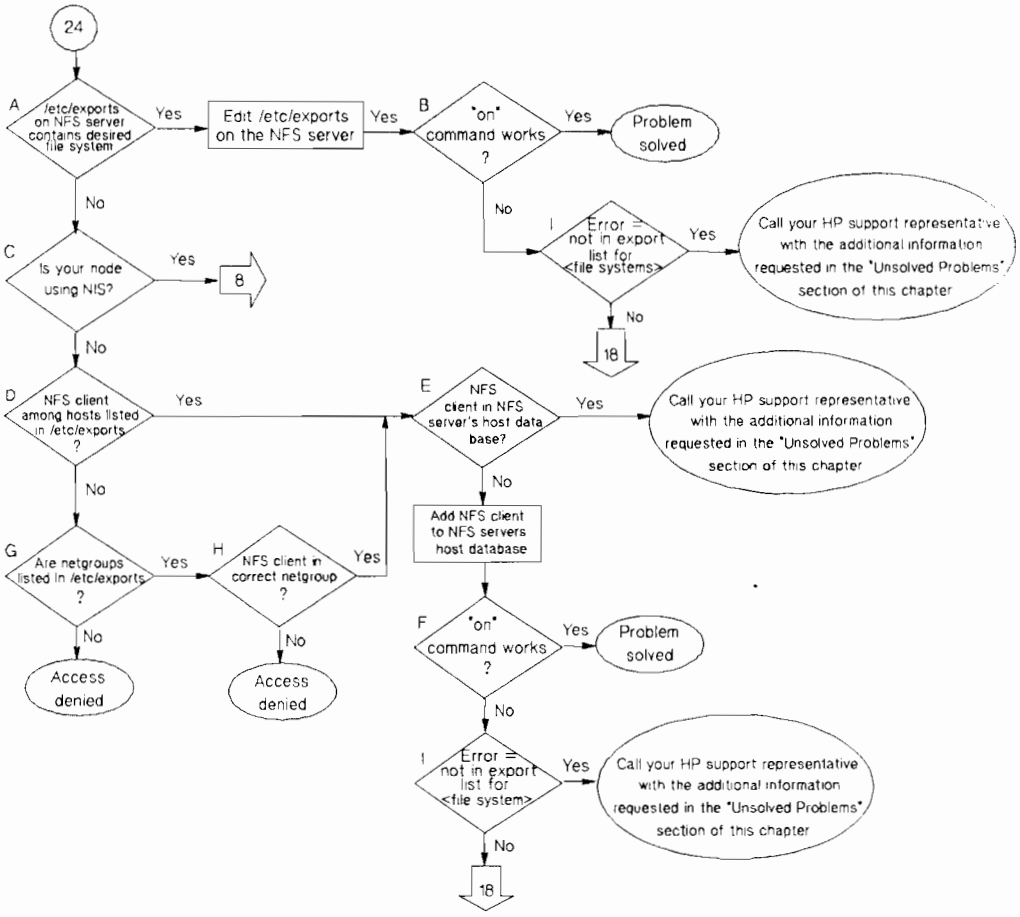

Flowchart 24: REX Server Denied Access through /etc/exports

REX Server Denied Access through /etc/exports (Flowchart 24)

Question	Yes: Action	No: Action
A. Does `/etc/exports` on the NFS server contain desired file system?	Edit `/etc/exports` on the NFS server to contain the file system rather that the directory. See B.	See C.
B. `on` command works now?	Problem solved.	See I.
C. Is your node using the Network Information Service (NIS)?	See Flowchart 8.	See D.
D. If hosts are listed in the desired `/etc/exports` entry, is the NFS client one of them?	See E.	See G.
E. Is the NFS client in the NFS server's host database?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Add NFS client to the NFS server's host database. See F.
F. `on` command works now?	Problem solved.	See I.
G. Are netgroups found for the desired `/etc/exports` entry?	See H.	Access for this client is denied.
H. Is the client included in a netgroup which is listed in the desired `/etc/exports` entry?	See E.	Access for this client is denied.

Question	Yes: Action	No: Action
I. Does the following error message appear? `on <server>: rexd: not in export list for file system`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

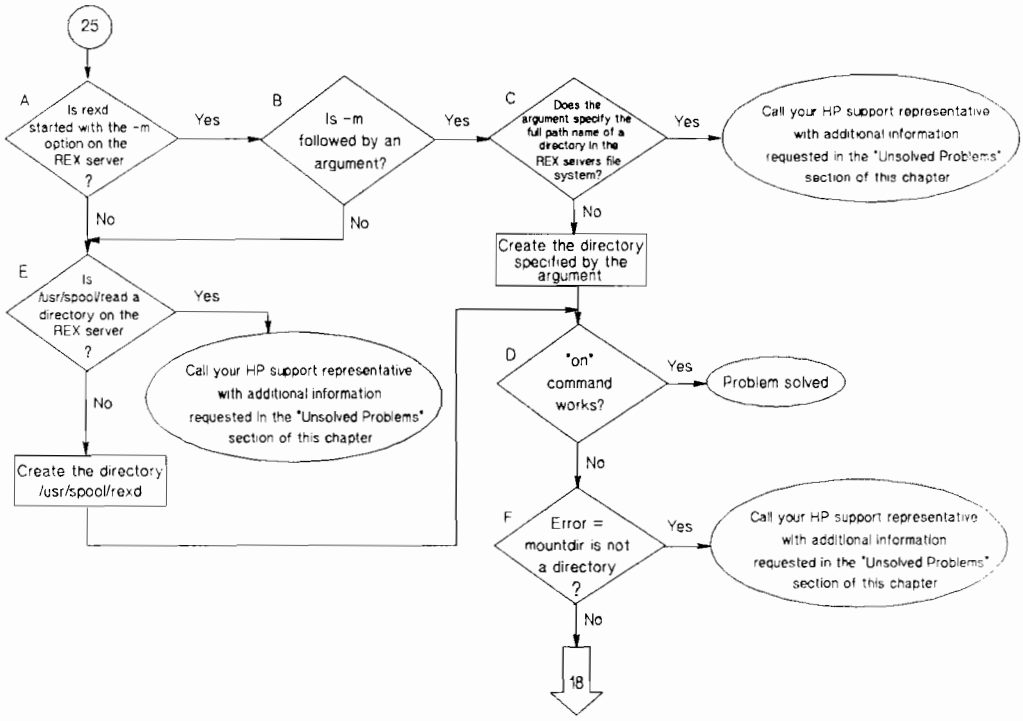

Flowchart 25: Mount Point Not a Directory

Mount Point Not a Directory (Flowchart 25)

Question	Yes: Action	No: Action
A. Is `rexd` on the REX server started with the `-m` option?	See B.	See E.
B. Is `-m` followed by a full path name?	See C.	See E.
C. Does the full path name specify a directory on the REX server?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Create the directory specified by the path name. See D.
D. `on` command works now?	Problem solved.	See F.
E. Is `/usr/spool/rexd` a directory on the REX server?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Create the directory `/usr/spool/rexd`. See D.
F. Does the following error message appear? `on <server>:rexd: (<mountdir>) is not a directory`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

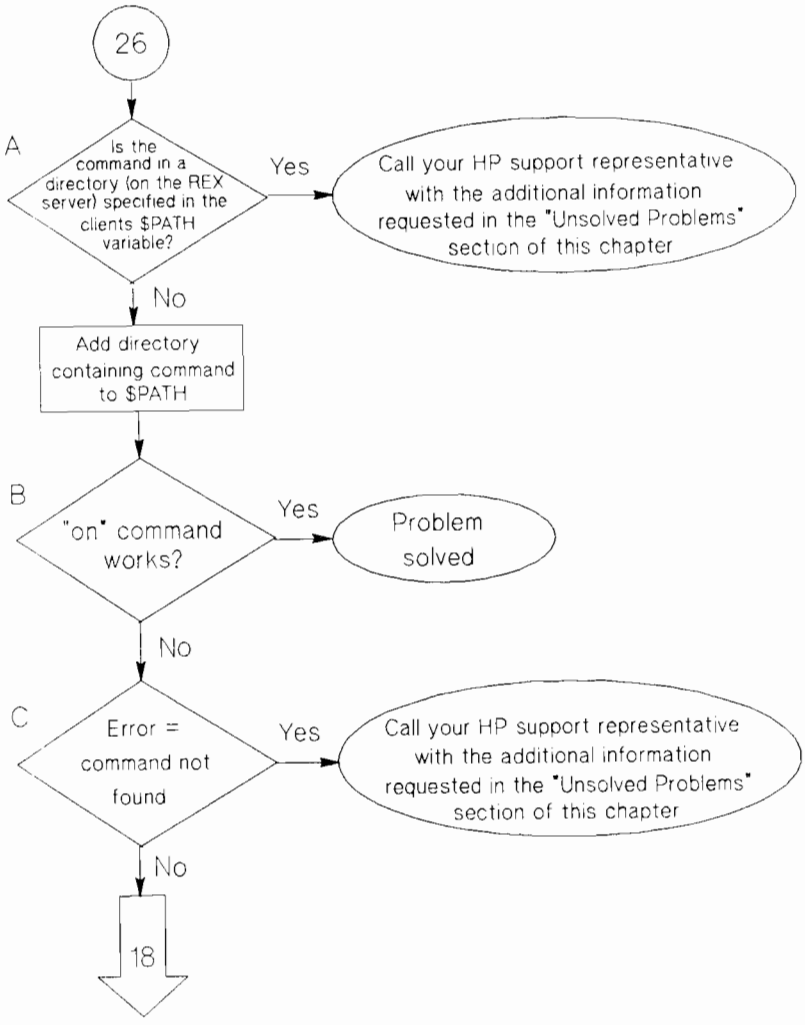

Flowchart 26: Command Not Found

Command Not Found (Flowchart 26)

Question	Yes: Action	No: Action
A. Is the command in a directory (visible on the REX server) which is specified in the user's `$PATH` variable?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Add the directory containing the command to the user's `$PATH` variable. See B.
B. `on` command works now?	Problem solved.	See C.
C. Does the following error message appear? `on: <server>:rexd: command not found`	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	See Flowchart 18.

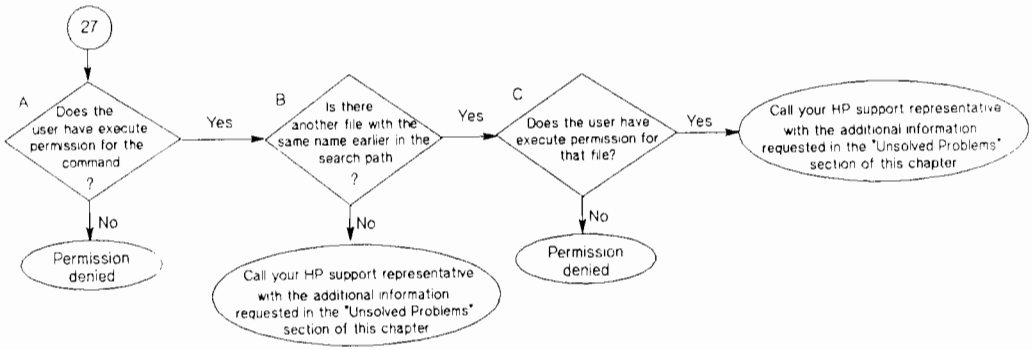

Flowchart 27: Permission Denied

Permission Denied (Flowchart 27)

Question	Yes: Action	No: Action
A. Does the user have execute permission for the command?	See B.	Permission denied.
B. Is there another file with the same name in a directory earlier in the user's $PATH variable?	See C.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
C. Does the user have execute permission for that file?	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.	Permission denied.

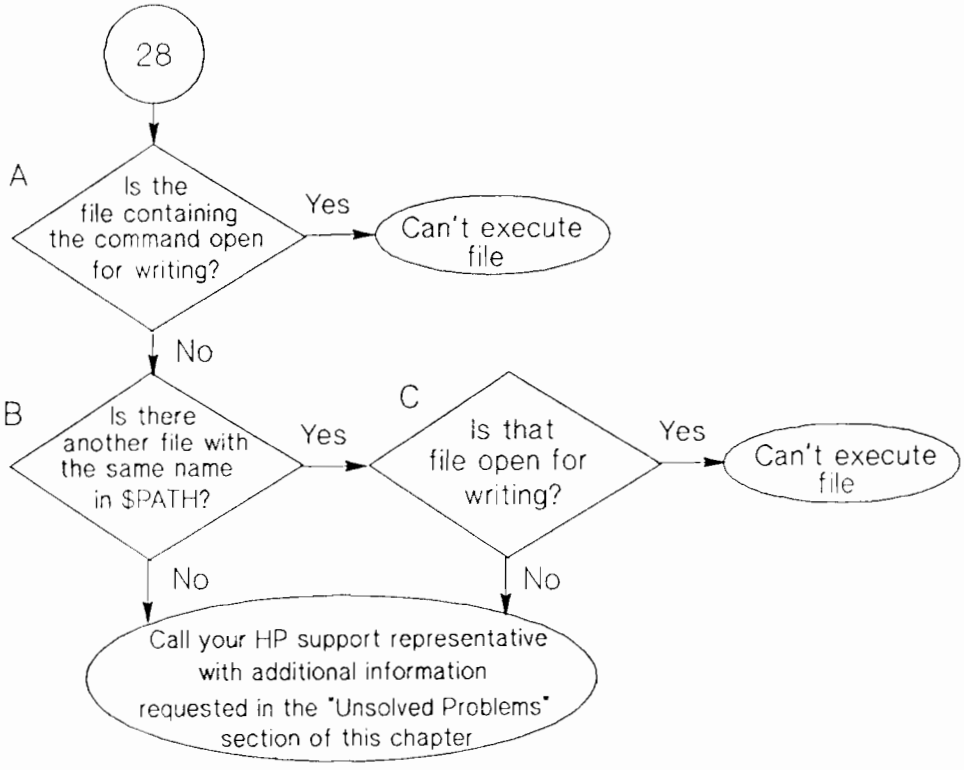

Flowchart 28: Text File Busy

Text File Busy (Flowchart 28)

Question	Yes: Action	No: Action
A. Is the file containing the command open for writing?	Can't execute file.	See B.
B. Is there another file with the same name in a directory earlier in the user's $PATH variable?	See C.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
C. Is that file currently open for writing?	Can't execute file.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

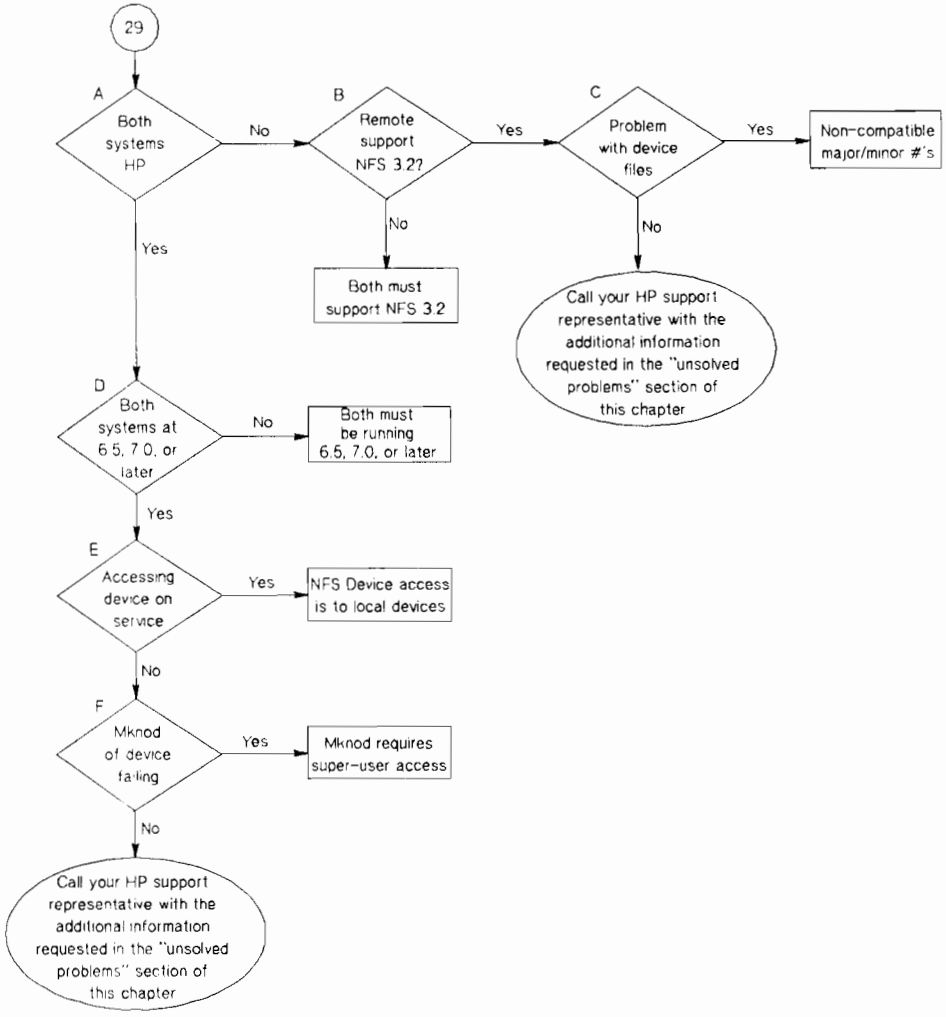

Flowchart 29: Device files/named pipes

Device files/named pipes (Flowchart 29)

Question	Yes: Action	No: Action
A. Are both systems HP systems?	See D.	See B.
B. Does the remote system support NFS 3.2 device files?	See C.	Both systems must support NFS 3.2. Consider mounting with the `-nodevs` option.
C. Is the problem with accessing device files?	Non-HP systems will have incompatible device major and minor numbers and format. This access will not work.	Problem is with named pipes. Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.
D. Are both systems running HP-UX releases 6.5, 7.0, or later?	See E.	Both systems must be running HP-UX release 6.5, 7.0, or later. Consider upgrading to the latest release of HP-UX.
E. Is the physical device represented by the device file attached to the server?	NFS device file access is to devices local to the client system.	See F.
F. Is an attempt to `mknod` a device file failing?	Creation of device files requires superuser access. Login as root on the server to create the device file.	Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.

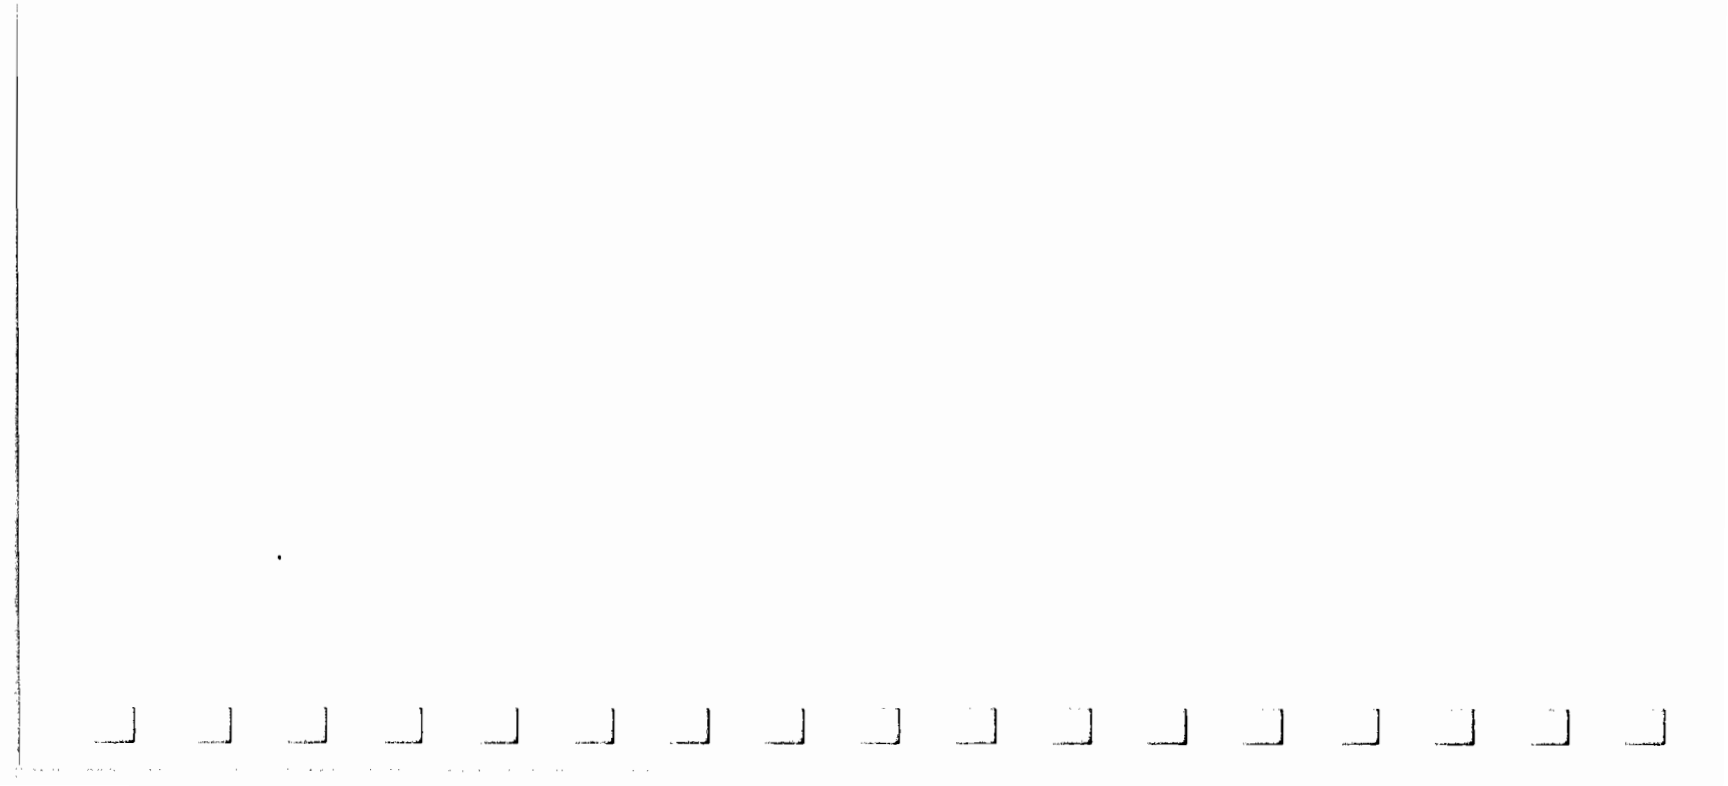

HP NFS Services vs. Local HP-UX

If you have applications running on HP-UX, they may behave differently over NFS Services. Use this appendix to understand the basic differences between NFS Services and local HP-UX operations.

HP NFS Services Networking Operation	Local HP-UX Operation
Append Mode	
If two processes operating on different clients open the same file using `O_APPEND`, the `write` operation may not append data to the file.	If two processes open the same file using `O_APPEND`, the `write` operation should append information to the file.
chacl(1)	
You can only use the `-F` option. The other options of `chacl` are not supported over NFS.	You can use all options locally.
Device Files	
NFS does not support remote access to device files, but does support local access to device files via NFS.	HP-UX supports local access to device files.
File Locking	
NFS supports remote file locking for NFS reads and writes in advisory mode only.	HP-UX supports local file locking in advisory and enforcement modes.

HP NFS Services Networking Operation	Local HP-UX Operation
getacl(2) system call	
Is not supported over NFS.	Is supported locally.
Group Membership	
A user may be a member of eight groups. If a user who is a member of more than eight groups attempts to access a file, the system accesses only the first eight groups for permission checking.	A user may be a member of up to 20 groups.
lseek(2)	
If two processes operating on different clients write to the same file, `lseek` with `whence=SEEK_END` may not set the file pointer to the desired location.	If two processes write to the same file, `lseek` with `whence=SEEK_END` should set the file pointer to the desired location.
mknod(1M) Command	
The `mknod` command will work only with named pipes over NFS.	You can use the `mknod` command locally for all file types.
Mount Points	
When operating in an HP-UX cluster environment, only file systems mounted on the cluster root server can contain mount points for NFS mounts. File systems mounted on cluster auxiliary servers cannot contain NFS mount points.	NFS mount points can exist on any mounted file system.
Named Pipes	
NFS named pipes cannot be used to communicate between machines in the same diskless cluster.	Named pipes can be used to communicate among clients in a diskless cluster.

HP NFS Services Networking Operation	Local HP-UX Operation
Reading Directories	
You cannot use the `read` call to read a remote directory, rather you should use `readdir`.	You can use the `read` call to read a local directory. However, to do so can restrict migration of programs to future HP-UX versions.
setacl(2) system call	
Is not supported over NFS.	Is supported locally.
setaclentry(3) library routine	
Is not supported over NFS.	Is supported locally.
Superuser Permission	
The superuser UID `0` is mapped to `-2` by default. Anything requiring superuser permission may not work over NFS. For example, a superuser may not be able to perform the following tasks: - Link and unlink directories. - Alter directories such as `/`, `/etc`, and `/bin`. - Use `chmod` to set sticky or setuid bits. - Do a `mknod` of device files.	Superuser has permission to perform any operation locally (by definition).

HP NFS Services Networking Operation	Local HP-UX Operation
System Time	
Commands that access clocks on different systems may not provide consistent times since system clocks differ. For example, if you give the `utime` command a `NULL` pointer for the times value, the following process occurs: 1. The system sets the access time and modification time according to the client node clock. 2. It then sends these times over to the server which changes the inode to reflect the new access and modification times. 3. The server node identifies the change in the inode and thus, modifies the inode's status change time according to its own clock. The result is a high probability of differing times between the server's access and modification times versus its status change time. *Note*: If operating in an HP-UX cluster environment, all nodes in the cluster have the same time as the root server's clock. Therefore, clock skew problems exist only if the root server's clock is different from other NFS servers.	Commands that access clocks on the local system provide consistent times.

HP NFS Services Networking Operation	Local HP-UX Operation
Unlinking	
The server does not keep state information and does not know if a process has a file open. See the following explanation: - The server will unlink a file if it receives a request to do so; thus, subsequent requests for the file will result in an error. - If a process opens a file and then unlinks it, the client renames the file so it appears to be gone. When the process quits, the client then unlinks the renamed file. - If the unlink request comes from a different node than from where the open request came from, the file is deleted.	If you open a local file and unlink it before you close the file, the file descriptor for the open file will still be valid to access the file.
yppasswd(1) Command vs. passwd(1) Command	
This command does not have a *password aging* feature. The superuser must know the current password to change another user's password. The password must contain: - At least five characters if it includes special characters and any combination of lowercase letters, uppercase letters, and numbers. Also, any combination that includes numbers, lowercase letters, and uppercase letters.	This command has a *password aging* feature. Superuser does not have to know the password to change another user's password. The following rules apply to the password: - Each password must have six or more characters: at least two alpha characters and at least one numeric or special character.

HP NFS Services Networking Operation	Local HP-UX Operation
yppasswd(1) Command vs. passwd(1) Command *(continued)*	
- At least six characters if it includes lowercase letters and numbers, uppercase letters and numbers, or a combination of lowercase and uppercase letters. - At least seven characters if it includes all lowercase letters or all uppercase letters.	- Each password must differ from the user's login name and any reverse or circular shift of that name. - New passwords must differ from the old by at least three characters.
pathconf/fpathconf	
The following variables for the `pathconf/fpathconf` system calls are not supported over NFS: _PC_CHOWN_RESTRICTED variable _PC_LINK_MAX variable _PC_NAME_MAX variable _PC_NO_TRUNC variable _PC_PATH_MAX variable The following variables for the `pathconf/fpathconf` system calls return local information over NFS: _PC_MAX_CANON variable _PC_MAX_INPUT variable _PC_VDISABLE variable The following variable for the `pathconf/fpathconf` systems calls is supported over NFS: _PC_PIPE_BUF variable	All variables are supported locally for the `pathconf/fpathconf` system calls: _PC_CHOWN_RESTRICTED variable _PC_LINK_MAX variable _PC_NAME_MAX variable _PC_NO_TRUNC variable _PC_PATH_MAX variable _PC_MAX_CANON variable _PC_MAX_INPUT variable _PC_VDISABLE variable _PC_PIPE_BUF variable

Moving from RFA to NFS

Remote File Access (RFA), one of the Network Services, has been discontinued. When you used networks consisting of all HP systems, RFA provided distributed file access among HP 9000 computers. In order to maintain distributed file access, you must move to NFS Services.

Why Move to NFS Services?

Using NFS Services in place of the RFA service has several advantages:

- NFS works with other vendors' equipment and other operating systems.
- NFS is a defacto industry standard.
- NFS allows transparent file access.
- NFS with the Network Information Service (NIS) provides centrally administered databases.

Use this appendix to translate your RFA applications to NFS applications.

Similarities

HP NFS Services and RFA have the following similarities:

- No remote device access.
- Not all UNIX® semantics are fully supported.

Differences

Refer to the following table for a list of differences between HP NFS and RFA.

NFS Services	RFA (Discontinued)
You can run `setuid` programs accessing data on remote file systems.	You cannot run `setuid` programs accessing data on remote file systems.
NFS operates in a heterogeneous operating system environment.	RFA operates on HP-UX operating systems only.
Only the superuser can perform remote NFS mounts.	All users can establish access to remote file systems.
You can centrally administer your databases using NIS.	You have no centrally administered database.
All users with read access to the mount point can read the remote file system.	Only users performing `netunam` can access the remote file systems.
Read and write file caching occurs on the clients; read caching occurs on the servers.	Read and write file caching occurs on the servers; caching does not occur on the clients.
The servers are stateless (do not remember client activities) and therefore, can be rebooted without interfering with client activities. (The client can resume access to the server when it is rebooted.)	The servers have state and therefore, remember the activities in which the client is involved.
One mount gives you access to only one file system.	One `netunam` gives you access to all file systems under the root directory.

Changing Scripts from RFA to NFS

Changing RFA scripts to NFS requires only minor changes. You can change both shell scripts that accept different path names and those that use hard-coded path names.

Shell Scripts that Accept Different Paths

Shell scripts that accept different paths require the following modifications:

- You must perform a remote mount of a file system or directory in *one* of the following ways:
  - As part of the script.
  - Before executing the script.

  Since superuser must execute mounts, the script must be `setuid root` if the mount is performed as part of the script.

Caution Having setuid root scripts is a potential security problem.

  If the script's owner does not have superuser permissions, the superuser can configure `/etc/checklist` to automatically mount the remote file systems at boot time. This process allows users to execute scripts without checking to see if the remote file system is accessible.

- Remove all calls to `netunam` from the script. Removing these calls prevents `netunam` failures from causing the scripts to fail.

Shell Scripts with Hard-Coded Paths

You can handle shell scripts with hard-coded path names in two ways:

- Change the path name in the script to correspond to the NFS mount point.
- Create a path name for the NFS mount point which corresponds to the path name in the script.

To mount the remote file system either as part of the script or automatically via `/etc/checklist`, you must modify the shell scripts as described in the previous section, "Shell Scripts that Accept Different Paths."

Change Pathnames

Change the path name in the script to correspond to the NFS mount point.

EXAMPLE: The script has a hard-coded path name of `/net/systemB/project`. Mount the remote directory `/project` on `/user/project` as follows:

```
mount systemB:/project /user/project
```

Now change the script to use the path name `/user/project` in place of `/net/systemB/project`.

Create New Pathnames

Create a path name for the NFS mount point that corresponds to the path name in the script.

EXAMPLE: The script has a hard-coded path name of `/net/systemB/project` which accesses the remote directory `/project`. To keep the path name the same:

1. Remove the network special file `/net/systemB`.

2. Create the directories `/net/systemB` and `/net/systemB/project`:

```
mount systemB:/project /net/systemB/project
```

Note

For RFA, access to the remote system occurred via a network special file. Creating an NFS mount point with the same name as the network special file for the remote system could cause confusion. Problems will not occur if you remove the network special file.

All remote access will then be via mount points that have the same names as the network special files that were removed.

NFS in an HP-UX Cluster Environment

Reference this appendix for interactions between NFS Services and HP-UX cluster environments using diskless capabilities.

HP-UX Cluster Terms

Term	Definition
Context Dependent File (CDF)	A hidden directory which contains all the versions of a file or directory needed by the different cnodes.
Cluster	One or more workstations linked together with a local area network (LAN), but consisting of only one root file system. For more information on cluster concepts, see *Managing Clusters of HP9000 Computers: Sharing the HP-UX Filing System*.
Cluster Auxiliary Server	A cluster client with a disk drive that contains files shared by the other members of the cluster.
Cluster Client	A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have its root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.
Cluster Node (Cnode)	Any node operating in an HP-UX cluster environment, including cluster clients and cluster servers.
Cluster Root Server	The only node in an HP-UX cluster that has the root file system directly attached to it.
Homogeneous Cluster	A diskless cluster composed of nodes of only one computer architecture (e.g., Series 300/400 only).
Mixed Cluster	Diskless cluster consisting of cnodes of multiple architectures.

NFS Configuration and Maintenance

Configure

If you configure NFS on the cluster root server, you must also configure NFS on all clients in the cluster. If the cluster root server does not have NFS configured, then none of the clients can use NFS.

Daemons

- The `nfsd` daemon should be running on the cluster root server and all cluster auxiliary servers if it is servicing NFS requests. Any `nfsd` daemon running on any other cnode is ignored.
- The `rpc.mountd` daemon should be running on the cluster root server and all cluster auxiliary servers if servicing NFS requests. Any `rpc.mountd` daemon running on any other cnode is ignored.
- The `biod` daemon should be running on all cnodes in the cluster.

Mount/Unmount

- If a cnode mounts a remote file system, all cnodes in the cluster can access the remote file system.
- If using NFS to mount a file system attached to a cluster, you must use the host name of the file system server as the node name specified in the `mount` command.
- If a cnode mounts a remote file system, any cnode in that cluster can unmount the remote file system.
- If a cnode unmounts a file system, all cnodes in the cluster will have that file system unmounted.
- Clients should not execute `umount -a`.
- NFS mount points may not exist on file systems mounted on cluster auxiliary servers.

Context Dependent Files (CDF)

When accessing a *CDF* via an NFS mount, the *CDF* member is chosen based on the context of the NFS server, not the accessing node. Since this access method may return unexpected results, HP recommends you do not mix *CDF*s with NFS.

Clock Skew

All nodes in the HP-UX cluster have the same time as the cluster root server's clock. Therefore, clock skew problems exists only if the cluster root server's clock is different from other NFS servers.

NIS Configuration and Maintenance

HP recommends that you execute `ypserv` only on the cluster root server for the two following reasons:

- For better performance.
- For assurance that the cluster root server is the only Network Information Service (NIS) server for that cluster.

Troubleshooting

You can troubleshoot NFS specific problems from the cluster root server and cluster auxiliary server as follows:

- If you are trying to mount an NFS file system, ensure you are using the cluster root server's host name as the node specified in the `mount` command.
- If problems exist in the link, cnodes will not be able to boot. Since link diagnostics reside on the root disk, first test the link from the cluster root server. (Refer to the *Installing and Administering LAN* manual.)

Password Security

This appendix explains the restrictions and limitations on the use of encrypted passwords and the secure password file with the Network Information Service (NIS). If you wish to review the normal use of passwords with the Network Information Service, see the "NIS Configuration and Maintenance" section in this manual. If you require additional information on the secure password file, see `passwd` in the *HP-UX Reference* manual.

The HP 9000 now supports a secure password file (`/.secure/etc/passwd`) used to hide your encrypted passwords from non-privileged users. Therefore, it is probable that if you use the secure password file, your `/etc/passwd` file will probably contain (in the password field) a character that is not part of the set of characters used in an encrypted password (e.g. *). The NIS database will not contain encrypted passwords if you use this `/etc/passwd` file to build your NIS password database. This prevents non-privileged users from reading your passwords, because anyone with access to NIS commands such as `ypcat` or NIS library routines such as `yp_first` and `yp_next` can read the NIS database.

If you are using the secure password file only to use the auditing subsystem and you do not need to hide your encrypted passwords, you can maintain an `/etc/passwd` file that contains encrypted passwords that match those in your secure password file. You can then use this `/etc/passwd` file to build your NIS database.

Note A password in the `/.secure/etc/passwd` file takes precedence over the password stored in NIS.

If you wish to hide the encrypted passwords in your HP systems and wish to continue to use the NIS password database to maintain other information kept on the password file, you can do the following:

- Build your NIS password database on the HP NIS master server using a password file that does not contain encrypted passwords (e.g. uses "*" in the password field).
- On an HP NIS client, maintain a copy of the secure password file so the passwords in that file will be used at login.
- On an HP or non-HP NIS client, maintain the encrypted password in the `/etc/passwd` file through an NIS escape.

EXAMPLE:

```
+username:encrypted_passwd:::::
```

E

Relinking Applications with RPC Library Functions

In the HP-UX 8.0 release, some changes were made to standard RPC library functions. Due to the changes, it may or may not be necessary to relink applications that call these functions. This appendix describes the changes that were made to RPC library functions and the impact of the changes on applications.

UDP Default Packet Size

In previous releases of HP-UX, the default maximum size of a UDP message was 9216 bytes. For HP-UX 8.0, this default changed to 2048 bytes. Prior to HP-UX 8.0, RPC calls used the default maximum UDP message size in many cases. Therefore, some applications linked on HP-UX systems prior to 8.0 will not use the 9216 byte default as before; but will use the 2048 byte default. If it is important that your RPC application has the ability to send and receive UDP messages of greater than 2048 bytes, you must relink your application.

RPC Broadcast

In previous releases of HP-UX, a process only needed superuser privileges to do a UDP broadcast. In HP-UX 8.0, any user may do a UDP broadcast as long as the UDP socket options are set correctly. The RPC call `clnt_broadcast()` did not set the socket options in previous releases. So for HP-UX 8.0, any application that relies on `clnt_broadcast()` must be relinked.

Maximum Number of File Descriptors

In previous releases of HP-UX, the maximum number of file descriptors that a single process could open at one time was 60. This limit has been changed for HP-UX 8.0. In HP-UX 8.0, the default limit is 128 file descriptors. This limit can be increased on a per-process basis using the `setrlimit` call. (See *setrlimit*(2) in the *HP-UX Reference*.) `setrlimit` allows a process to increase the maximum number of open file descriptors up to 1024 for non-superuser processes and 2048 for superuser processes.

Due to this change in HP-UX, the RPC calls are required to handle `select` calls with large file descriptor numbers. (See *select*(2) in the *HP-UX Reference*.) Relinking your application as a result of this change is only necessary if the application is spawned by a process that has increased the maximum number of open file descriptors allowed (via the `setrlimit` call).

Glossary

Alias

A term for referencing alternate networks, hosts, and protocols names.

ARPA

Advanced Research Projects Agency

A U.S. government agency that was instrumental in developing and using the original ARPA Services networking standards.

Bind

A process by which a client locates and directs all requests for data to a specific server.

A process of establishing the address of a socket that allows other sockets to connect to it or to send data to it.

An acronym for Berkeley Internet Name Domain. The BIND Name Server is a distributed network lookup service.

CDF (Context Dependent File)

A hidden directory that contains all the versions of a file needed by the different cnodes.

Client

A node that requests data or services from other nodes (servers).

A process that requests other processes to perform operations.

Note: An NFS client can also be configured as any combination of an NFS server, NIS client, or NIS server. (An NIS server must also be configured as an NIS client.)

Clock Skew

A difference in clock times between systems.

Cluster

One or more workstations linked together with a local area network (LAN), but consisting of only one root file system.

Cluster Auxiliary Server

A cluster client with a disk drive that contains files shared by the other members of the cluster.

Cluster Client

A node in an HP-UX cluster that uses networking capabilities to share file systems,, but does not have its root file system directly attached. For HP-UX 8.0, cluster clients can have locally mounted disks for local data storage.

Cluster Root Server

The only node in an HP-UX cluster that has the root file system directly attached to it.

Cnode (Cluster Node)

Any node operating in an HP-UX cluster environment, including cluster clients, cluster auxiliary servers, and the cluster root server.

Daemon

Background programs that are always running, waiting for a request to perform a task.

Escape Sequence (NIS)

Characters used within files to force inclusion and exclusion of data from NIS databases. The escape sequences are as follows.

* + (plus sign)
* - (minus sign)
* +*@netgroup_name*
* -*@netgroup_name*

Export

To make a file system available to remote nodes via NFS.

External Data Representation (XDR)

A protocol that translates machine-dependent data formats (i.e., internal representations) to a universal format used by other network hosts using XDR.

File System

A directory structure used to organize files.

GID

A value that identifies a group in HP-UX.

Global (NIS)

A means of access in which the system always reads NIS maps rather than the local ASCII files.

Hard Mount

A mount that causes NFS to retry a remote file system request until it succeeds, you interrupt it (default option), or you reboot the system.

Home Node

A term used in Virtual Home Environment (VHE) to refer to the machine on which a user's home directory physically resides.

Host

A node that has primary functions other than switching data for the network.

Host Node

A term used in Virtual Home Environment (VHE) to refer to the node a user is logged in to. This node environment is set up from the configuration files found on the user's home node.

Import

To obtain access to a remote file system from an outside source; to mount.

Internet Address

A four-byte quantity that is distinct from a link-level address and is the network address of a computer node. This address identifies both the specific network and the specific host on the network.

Interruptable Mount

A mount that allows you to interrupt an NFS request by pressing an interrupt key. (Though the interrupt key is not standardized, common ones include [CTRL] - [C] and [BREAK].)

Key (NIS)

A string of characters (no imbedded blanks or tabs) that indexes the values within an NIS map so the system can easily retrieve information. For example, in the `passwd.byname` map, the users' login names are the keys and the matching lines from `/etc/passwd` are the values.

Local (NIS)

A means of access in which the system first reads the local ASCII file. If it encounters an escape sequence, it then accesses the NIS databases.

Map (NIS)

A file consisting of logical records; a search key and related value form each record. NIS clients can request the value associated with any key within a map.

NIS map is synonymous with NIS database.

Map Nickname (NIS)

A synonym for the NIS map name when using certain NIS commands.

Master Server (NIS)

The node on which one or more NIS maps are constructed from ASCII files. These maps are then copied to the NIS slave servers for the NIS clients to access.

Mount

To obtain access to a remote or local file system or directory (import).

Mount Point

The name of the directory on which a file system is mounted.

Netgroup

A network-wide group of nodes and users defined in `/etc/netgroup`.

Network Information Service (NIS)

An optional network service composed of databases (maps) and processes that provide NIS clients access to the maps. The NIS service enables you to administer these databases from one node.

NIS may or may not be active; check with your system administrator.

Network Lock Manager

A facility for locking files and synchronizing access to shared files.

Network Status Monitor

A daemon running on all network computers to maintain stateful locking service within NFS. It also allows applications to monitor the status of other computers.

NFS

Network File System.

NIS Client

A node that requests data or services from NIS servers.

An NIS process that requests other NIS processes to perform operations.

Note: An NIS client can also be configured as any combination of an NIS server, NFS client, or NFS server. (An NIS server must also be configured as an NIS client.)

NIS Database

See Map (NIS).

NIS Domain

A logical grouping of NIS maps (databases) stored in one location. NIS domains are specific to the NIS network service and are not associated with other network domains.

NIS Map

See Map (NIS).

NIS Password

The password for a user's login ID that exists in the NIS `passwd` map. The NIS password is the same one as the user password, but is administered through the NIS.

You do not have to have an NIS password to access the NIS databases.

NIS Server

A node that provides data (maps) or services to other nodes (NIS clients) on the network using NIS.

An NIS process that performs operations as requested by other NIS processes.

Note: An NIS server must also be configured as an NIS client. It can also be configured as an NFS server, NFS client, or both.

Node

A computer system that is attached to or is part of a computer network.

Propagate

To copy maps (data) from one NIS server to another.

Protocol

The rules and steps by which servers and clients exchange data and control information.

Remote Execution Facility (REX)

A facility which allows a user to execute commands on a remote node.

Remote Procedure Call (RPC)

A call made by clients either to access server information or to request action from servers.

Remote Procedure Call Protocol Compiler (RPCGEN)

A remote procedure call compiler used to help programmers write RPC applications by automatically generating necessary programs and code fragments.

Server

A node that provides data or services to other nodes (clients) on the network.

A process that performs operations as requested by other processes.

Note: An NFS server can also be configured as any combination of an NFS client, NIS client, or NIS server. (An NIS server must also be configured as an NIS client.)

Slave Server (NIS)

A node that copies NIS maps from the NIS master server and then provides NIS clients access to these maps.

Soft Mount

An optional mount that causes access to remote file systems to abort requests after one NFS attempt.

Stateless Server

Servers do not maintain (preserve) information relating to each file being served. Each file request moves across the network with the parameters attached to it locally (e.g., read and write privileges).

Steady State

Servers maintain (preserve) information relating to each file being served.

For NIS, the information contained in an NIS map is consistent among all NIS servers within a given NIS domain (i.e., is not in the process of being updated).

UID

A value that identifies a user in HP-UX.

Unmount

To remove access rights to a file system or disk that was mounted via the `mount` command.

update

The HP-UX command that installs software onto the system.

Value (NIS)

A unit of information stored in NIS maps; each value has a corresponding key (index) so the system can easily retrieve it. For example, in the `passwd.byname` map, the users' login names are the keys and the matching lines from `/etc/passwd` are the values.

VHE

See Virtual Home Environment.

Virtual Home Environment (VHE)

A network service that allows users to log in at host nodes and utilize their home nodes' execution environments.

XDR

See External Data Representation.

Index

D

E

F

G

H

I

K

L

M

N

P

R

S

T

U

V

W

X

Y

SALES & SUPPORT OFFICES
Arranged alphabetically by country

Product Line Sales/Support Key

Key Product Line
A Analytical
CM Components
C Computer Systems
E Electronic Instruments & Measurement Systems
M Medical Products
P Personal Computation Products
*** Sales only for specific product line**
**** Support only for specific product line**

IMPORTANT: These symbols designate general product line capability. They do not insure sales or support availability for all products within a line, at all locations. Contact your local sales office for information regarding locations where HP support is available for specific products.

HEADQUARTERS OFFICES

If there is no sales office listed for your area, contact one of these headquarters offices.

ASIA
Hewlett-Packard Asia Ltd.
47/F, 26 Harbour Rd.,
Wanchai, **HONG KONG**
G.P.O. Box 863, Hong Kong
Tel: 5-8330833
Telex: 76793 HPA HX
Cable: HPASIAL TD

CANADA
Hewlett-Packard (Canada) Ltd.
6877 Goreway Drive
MISSISSAUGA, Ontario L4V 1M8
Tel: (416) 678-9430
Telex: 069-8644

EASTERN EUROPE
Hewlett-Packard Ges.m.b.h.
Lieblgasse 1
P.O.Box 72
A-1222 **VIENNA**, Austria
Tel: (222) 2500-0
Telex: 1 3 4425 HEPA A

NORTHERN EUROPE
Hewlett-Packard S.A.
V. D. Hooplaan 241
P.O.Box 999
NL-118 LN 15 **AMSTELVEEN**
The Netherlands
Tel: 20 5479999
Telex: 18919 hpner

SOUTH EAST EUROPE
Hewlett-Packard S.A.
World Trade Center
110 Avenue Louis-Casai
1215 Cointrin, **GENEVA**, Switzerland
Tel: (022) 98 96 51
Telex: 27225 hpser
Mail Address
P.O. Box
CH-1217 Meyrin 1
GENEVA
Switzerland

MIDDLE EAST AND CENTRAL AFRICA
Hewlett-Packard S.A.
Middle East/Central
Africa Sales H.Q.
7, rue du Bois-du-Lan
P.O. Box 364
CH-1217 Meyrin 1
GENEVA
Switzerland
Tel: (022) 83 12 12
Telex: 27835 hmea ch
Telefax: (022) 83 15 35

UNITED KINGDOM
Hewlett-Packard Ltd.
Nine Mile Ride
WOKINGHAM
Berkshire, RG11 3LL
Tel: 0344 773100
Telex: 848805/848814/848912

UNITED STATES OF AMERICA
Customer Information Center
(800) 752-0900
6:00 AM to 5 PM Pacific Time

EASTERN USA
Hewlett-Packard Co.
4 Choke Cherry Road
ROCKVILLE, MD 20850
Tel: (301) 948-6370

MIDWESTERN USA
Hewlett-Packard Co.
5201 Tollview Drive
ROLLING MEADOWS, IL 60008
Tel: (312) 255-9800

SOUTHERN USA
Hewlett-Packard Co.
2000 South Park Place
ATLANTA, GA 30339
Tel: (404) 955-1500

WESTERN USA
Hewlett-Packard Co.
5161 Lankershim Blvd.
NORTH HOLLYWOOD, CA 91601
Tel: (818) 505-5600

OTHER INTERNATIONAL AREAS
Hewlett-Packard Co.
Intercontinental Headquarters
3495 Deer Creek Road
PALO ALTO, CA 94304
Tel: (415) 857-1501
Telex: 034-8300
Cable: HEWPACK

ALGERIA
Hewlett-Packard Trading S.A.
Bureau de Liaison Alger
Villa des Lions
9, rue Galloul
DZ-BORDJ EL BAHRI
Tel: 76 03 36
Telex: 63343 dhion dz

ANGOLA
Telectra Angola LDA
Empresa Técnica de Equipamentos
16 rue Cons. Julio de Vilhena
LUANDA
Tel: 35515,35516
Telex: 3134
E,P

ARGENTINA
Hewlett-Packard Argentina S.A.
Montaneses 2140/50
1428 **BUENOS AIRES**
Tel: 541-11-1441
Telex: 22796 HEW PAC-AR
A,C,E,P
Biotron S.A.C.I.M.e.I.
Av. Paso Colon 221, Piso 9
1399 **BUENOS AIRES**
Tel: 541-333-490,
541-322-587
Telex: 17595 BIONAR
M
Laboratorio Rodriguez
Corswant S.R.L.
Misiones, 1156 - 1876
Bernal, Oeste
BUENOS AIRES
Tel: 252-3958, 252-4991
A
Intermaco S.R.L.
Florida 537/71
Galeria Jardin - Local 28
1005 **BUENOS AIRES**
Tel: 393-4471/1928
Telex: 22796 HEW PAC-AR
P (Calculators)
Argentina Esanco S.R.L.
A/ASCO 2328
1416 **BUENOS AIRES**
Tel: 541-58-1981, 541-59-2767
Telex: 22796 HEW PAC-AR
A
All Computers S.A.
Montaneses 2140/50 5 Piso
1428 **BUENOS AIRES**
Tel: 781-4030/4039/783-4886
Telex: 18148 Ocme
P

AUSTRALIA

Adelaide, South Australia Office
Hewlett-Packard Australia Ltd.
153 Greenhill Road
PARKSIDE, S.A. 5063
Tel: 61-8-272-5911
Telex: 82536
Cable: HEWPARD Adelaide
A*,C,CM,E,P

Brisbane, Queensland Office
Hewlett-Packard Australia Ltd.
10 Payne Road
THE GAP, Queensland 4061
Tel: 61-7-300-4133
Telex: 42133
Cable: HEWPARD Brisbane
A,C,CM,E,M,P

Canberra, Australia Capital Territory Office
Hewlett-Packard Australia Ltd.
Thynne Street, Fern Hill Park
BRUCE, A.C.T. 2617
P.O. Box 257,
JAMISON, A.C.T. 2614
Tel: 61-62-80-4244
Telex: 62650
Cable: HEWPARD Canberra
C,CM,E,P

Melbourne, Victoria Office
Hewlett-Packard Australia Ltd.
31-41 Joseph Street
P.O. Box 221
BLACKBURN, Victoria 3130
Tel: 61-3-895-2895
Telex: 31-024
Cable: HEWPARD Melbourne
A,C,CM,E,M,P

Perth, Western Australia Office
Hewlett-Packard Australia Ltd.
Herdsman Business Park
CLAREMONT, W.A. 6010
Tel: 61-9-383-2188
Telex: 93859
Cable: HEWPARD Perth
C,CM,E,P

Sydney, New South Wales Office
Hewlett-Packard Australia Ltd.
17-23 Talavera Road
P.O. Box 308
NORTH RYDE, N.S.W. 2113
Tel: 61-2-888-4444
Telex: 21561
Cable: HEWPARD Sydney
A,C,CM,E,M,P

AUSTRIA
Hewlett-Packard Ges.m.b.h.
Verkaufsbuero Graz
Grottenhofstrasse 94
A-8052 **GRAZ**
Tel: 43-316-291-5660
Telex: 312375
C,E
Hewlett-Packard Ges.m.b.h.
Lieblgasse 1
P.O. Box 72
A-1222 **VIENNA**
Tel: 43-222-2500
Telex: 134425 HEPA A
A,C,CM,E,M,P

BAHRAIN
Green Salon
P.O. Box 557
MANAMA
Tel: 255503-250950
Telex: 84419
P
Wael Pharmacy
P.O. Box 648
MANAMA
Tel: 256123
Telex: 8550 WAEL BN
E,M
Zayani Computer Systems
218 Shaik Mubarak Building
Government Avenue
P.O. Box 5918
MANAMA
Tel: 276278
Telex: 9015 plans bn
P

BELGIUM
Hewlett-Packard Belgium S.A./N.V.
Blvd de la Woluwe, 100
Woluwedal
B-1200 **BRUSSELS**
Tel: (02) 32-2-761-31-11
Telex: 23494 hewpac
A,C,CM,E,M,P

BERMUDA
Applied Computer Technologies
Atlantic House Building
P.O. Box HM 2091
Par-La-Ville Road
HAMILTON 5
Tel: 295-1616
Telex: 380 3589/ACT BA
P

BOLIVIA
Arrellano Ltda.
Av. 20 de Octubre #2125
Casilla 1383
LA PAZ
Tel: 368541
M

BRAZIL
Hewlett-Packard do Brasil S.A.
Alameda Rio Negro, 750-I. AND.
ALPHAVILLE
06400 Barueri SP
Tel: (011) 421.1311
Telex: (011) 71351 HPBR BR
Cable: HEWPACK Sao Paulo
CM,E
Hewlett-Packard do Brasil S.A.
Praia de Botafago 228-A-614
6. AND.-CONJ. 601
Edificio Argentina - Ala A
22250 **RIO DE JANEIRO, RJ**
Tel: (021) 552-6422
Telex: 21905 HPBR BR
Cable: HEWPACK Rio de Janeiro
E
Van Den Ciantrica Ltda.
Rua Jose Bonifacio, 458
Todos os Santos
20771 **RIO DE JANEIRO, RJ**
Tel: (021) 593-8223
Telex: 33487 EGLB BR
A
ANAMED I.C.E.I. Ltda.
Rua Verguerio, 360
04012 **SAO PAULO, SP**
Tel: (011) 572-1106
Telex: 24720 HPBR BR
M
Datatronix Electronica Ltda.
Av. Pacaembu 746-C11
SAO PAULO, SP
Tel: (118) 260111
CM

BRUNEI

Komputer Wisman Sdn Bhd
G6, Chandrawaseh Cmplx.
Jalan Tutong
P O Box 1297
BANDAR SERI BEGAWAN
NEGARA BRUNI DARUSSALAM
Tel: 673-2-2000-70/26711
C,E,P

CAMEROON

Beriac
B P 23
DOUALA
Tel: 420153
Telex: 5351
C,P

CANADA

Alberta

Hewlett-Packard (Canada) Ltd
3030 3rd Avenue N E.
CALGARY, Alberta T2A 6T7
Tel: (403) 235-3100
A,C,CM,E*,M,P*

Hewlett-Packard (Canada) Ltd.
11120-178th Street
EDMONTON, Alberta T5S 1P2
Tel: (403) 486-6666
A,C,CM,E,M,P

British Columbia

Hewlett-Packard (Canada) Ltd
10691 Shellbridge Way
RICHMOND,
British Columbia V6X 2W8
Tel: (604) 270-2277
Telex: 610-922-5059
A,C,CM,E*,M,P*

Hewlett-Packard (Canada) Ltd
121 - 3350 Douglas Street
VICTORIA, British Columbia V8Z 3L1
Tel: (604) 381-6616
C

Manitoba

Hewlett-Packard (Canada) Ltd
1825 Inkster Blvd
WINNIPEG, Manitoba R2X 1R3
Tel: (204) 694-2777
A,C,CM,E,M,P*

New Brunswick

Hewlett-Packard (Canada) Ltd
814 Main Street
MONCTON, New Brunswick E1C 1E6
Tel: (506) 855-2841
C

Nova Scotia

Hewlett-Packard (Canada) Ltd
Suite 111
900 Windmill Road
DARTMOUTH, Nova Scotia B3B 1P7
Tel: (902) 469-7820
C,CM,E*,M,P*

Ontario

Hewlett-Packard (Canada) Ltd
3325 N. Service Rd., Unit W03
BURLINGTON, Ontario L7N 3G2
Tel: (416) 335-8644
C,M*

Hewlett-Packard (Canada) Ltd
552 Newbold Street
LONDON, Ontario N6E 2S5
Tel: (519) 686-9181
A,C,CM,E*,M,P*

Hewlett-Packard (Canada) Ltd
6877 Goreway Drive
MISSISSAUGA, Ontario L4V 1M8
Tel: (416) 678-9430
Telex: 069-83644
A,C,CM,E,M,P

Hewlett-Packard (Canada) Ltd.
2670 Queensview Dr.
OTTAWA, Ontario K2B 8K1
Tel: (613) 820-6483
A,C,CM,E*,M,P*

Hewlett-Packard (Canada) Ltd
3790 Victoria Park Ave
WILLOWDALE, Ontario M2H 3H7
Tel: (416) 499-2550
C,E

Quebec

Hewlett-Packard (Canada) Ltd.
17500 Trans Canada Highway
South Service Road
KIRKLAND, Quebec H9J 2X8
Tel: (514) 697-4232
Telex: 058-21521
A,C,CM,E,M,P*

Hewlett-Packard (Canada) Ltd.
1150 rue Claire Fontaine
QUEBEC CITY, Quebec G1R 5G4
Tel: (418) 648-0726
C

Hewlett-Packard (Canada) Ltd.
130 Robin Crescent
SASKATOON, Saskatchewan S7L 6M7
Tel: (306) 242-3702
C

CHILE

ASC Ltda
Austria 2041
SANTIAGO
Tel: 223-5946, 223-6148
Telex: 392-340192 ASC CK
C,P

Jorge Calcagni y Cia
Av. Italia 634 Santiago
Casilla 16475
SANTIAGO 9
Tel: 9-011-562-222-0222
Telex: 392440283 JCYCL CZ
CM,E,M

Metrolab S.A.
Monjitas 454 of. 206
SANTIAGO
Tel: 395752, 398296
Telex: 340866 METLAB CK
A

Olympia (Chile) Ltda.
Av. Rodrigo de Araya 1045
Casilla 256-V
SANTIAGO 21
Tel: 225-5044
Telex: 340892 OLYMP
Cable: Olympiachile Santiagochile
C,P

CHINA, People's Republic of

China Hewlett-Packard Co., Ltd
47/F China Resources Bldg.
26 Harbour Road
HONG KONG
Tel: 5-8330833
Telex: 76793 HPA HX
Cable: HP ASIA LTD
A*,M*

China Hewlett-Packard Co., Ltd
P.O. Box 9610, Beijing
4th Floor, 2nd Watch Factory Main
Shuang Yu Shou, Bei San Huan Road
Hai Dian District
BEIJING
Tel: 33-1947 33-7426
Telex: 22601 CTSHP CN
Cable: 1920 Beijing
A,C,CM,E,M,P

China Hewlett-Packard Co., Ltd
CHP Shanghai Branch
23/F Shanghai Union Building
100 Yan An Rd. East
SHANG-HAI
Tel: 265550
Telex: 33571 CHPSB CN
Cable: 3416 Shanghai
A,C,CM,E,M,P

COLOMBIA

Instrumentación
H. A. Langebaek & Kier S.A.
Carrerra 4A No. 52A-26
Apartado Aereo 6287
BOGOTA 1, D.E.
Tel: 212-1466
Telex: 44400 INST CO
Cable: AARIS Bogota
CM,E,M

Nefromedicas Ltda.
Calle 123 No. 9B-31
Apartado Aereo 100-958
BOGOTA D.E., 10
Tel: 213-5267, 213-1615
Telex: 43415 HEGAS CD
A

Compumundo
Avenida 15 # 107-80
BOGOTA D.E.
Tel: 57-214-4458
Telex: 39645466 MARCO
P

Carvajal, S.A.
Calle 29 Norte No. 6A-40
Apartado Aereo 46
CALI
Tel: 9-011-57-3-621888
Telex: 39655650 CUJCL CO
C,E,P

CONGO

Seric-Congo
B.P. 2105
BRAZZAVILLE
Tel: 815034
Telex: 5262

COSTA RICA

Cientifica Costarricense S.A.
Avenida 2, Calle 5
San Pedro de Montes de Oca
Apartado 10159
SAN JOSE
Tel: 9-011-506-243-820
Telex: 3032367 GALGUR CR
CM,E,M

O. Fischel R. Y. Cia. S.A.
Apartados 434-10174
SAN JOSE
Tel: 23-72-44
Telex: 2379
Cable: OFIR
A

CYPRUS

Telerexa Ltd
P.O. Box 1152
Valentine House
8 Stassandrou St.
NICOSIA
Tel: 45 628, 62 698
Telex: 5645 tirx cy
E,M,P

DENMARK

Hewlett-Packard A/S
Kongevejen 25
DK-3460 **BIRKEROD**
Tel: 45-02-81-6540
Telex: 37409 hpas dk
A,C,CM,E,M,P

Hewlett-Packard A/S
Rolighedsvej 32
DK-8240 **RISSKOV**, Aarhus
Tel: 45-06-17-6000
Telex: 37409 hpas dk
C,E

DOMINICAN REPUBLIC

Microprog S.A.
Juan Tomás Mejía y Cotes No. 60
Arroyo Hondo
SANTO DOMINGO
Tel: 565-6268
Telex: 4510 ARENTA DR (RCA)
P

ECUADOR

CYEDE Cia. Ltda.
Avenida Eloy Alfaro 1749
y Belgica
Casilla 6423 CCI
QUITO
Tel: 9-011-593-2-450975
Telex: 39322548 CYEDE ED
E,P

Medtronics
Valladolid 524 Madrid
P.O. 9171, **QUITO**
Tel: 2-238-951
Telex: 2298 ECUAME ED
A

Hospitalar S.A.
Robles 625
Casilla 3590
QUITO
Tel: 545-250, 545-122
Telex: 2485 HOSPTL ED
Cable: HOSPITALAR-Quito
M

Ecuador Overseas Agencies C.A.
Calle 9 de Octubre #818
P.O. Box 1296, Guayaquil
QUITO
Tel: 306022
Telex: 3361 PBCGYE ED
M

EGYPT

Sakrco Enterprises
P.O. Box 259
ALEXANDRIA
Tel: 802908, 808020, 805302
Telex: 54333
C

International Engineering Associates
6 El Gamea Street
Agouza
CAIRO
Tel: 71-21-68134-80-940
Telex: 93830 IEA UN
Cable: INTEGASSO
E

Sakrco Enterprises
70 Mossadak Street
Dokki, Giza
CAIRO
Tel: 706 440, 701 087
Telex: 9337
C

S.S.C. Medical
40 Gezeral El Arab Street
Mohandessin
CAIRO
Tel: 803844, 805998, 810263
Telex: 20503 SSC UN
M*

EL SALVADOR

IPESA de El Salvador S.A.
29 Avenida Norte 1223
SAN SALVADOR
Tel: 9-011-503-266-858
Telex: 301 20539 IPESA SAL
A,C,CM,E,P

ETHIOPIA

Seric-Ethiopia
P.O. Box 2764
ADDIS ABABA
Tel: 185114
Telex: 21150
C,P

FINLAND

Hewlett-Packard Finland
Field Oy
Nittytanpoiku IO
00620 **HELSINKI**
Tel: (90) 757-1011
Telex: 122022 Field SF
CM
Hewlett-Packard Oy
Piispankallionite 17
02200 **ESPOO**
Tel: (90) 887-21
Telex: 121563 HEWPA SF
A,C,E,M,P

FRANCE

Hewlett-Packard France
Z.I. Mercure B
Rue Berthelot
13763 Les Milles Cedex
AIX-EN-PROVENCE
Tel: 33-42-59-4102
Telex: 410770F
A,C,E,M

Hewlett-Packard France
64, Rue Marchand Saillant
F-61000 **ALENCON**
Tel: (33) 29 04 42
C**

Hewlett-Packard France
Batiment Levitan
2585, route de Grasse
Bretelle Autoroute
06600 **ANTIBES**
Tel: (93) 74-59-19
C

FRANCE (Cont'd)

Hewlett-Packard France
28 Rue de la République
Boite Postale 503
25026 **BESANÇON CEDEX, FRANCE**
Tel: (81) 83-16-22
Telex: 381157
C,E*

Hewlett-Packard France
ZA Kergaradec
Rue Fernand Forest
F-29239 **GOUESNOU**
Tel: (98) 41-87-90
E

Hewlett-Packard France
Chemin des Mouilles
Boite Postale 162
69131 **ECULLY** Cedex (Lyon)
Tel: 33-78-33-8125
Telex: 310617F
A,C,E,M,P*

Hewlett-Packard France
Parc d'activités du Bois Briard
2 Avenue du Lac
F-91040 **EVRY** Cedex
Tel: 3311/6077 9660
Telex: 892315F
C

Hewlett-Packard France
Application Center
5, avenue Raymond Chanas
38320 **EYBENS** (Grenoble)
Tel: (76) 62-57-98
Telex: 980124 HP GRENOB EYBE
C

Hewlett-Packard France
Rue Fernand. Forest
Z.A. Kergaradec
29239 **GOUESNOU**
Tel: (98) 41-87-90

Hewlett-Packard France
Parc Club des Tanneries
Batiment B4
4, Rue de la Farsanderie
67381 **LINCOLSHEIM**
(Strasbourg)
Tel. (88) 76-15-00
Telex: 890141F
C,E*,M*,P*

Hewlett-Packard France
Centre d'affaires Paris-Nord
Bâtiment Ampère
Rue de la Commune de Paris
Boite Postale 300
93153 **LE BLANC-MESNIL**
Tel: (1) 865-44-52
Telex: 211032F
C,E,M

Hewlett-Packard France
Parc d'activités Cadéra
Quartier Jean-Mermoz
Avenue du Président JF Kennedy
33700 **MERIGNAC** (Bordeaux)
Tel: 33-56-34-0084
Telex: 550105F
C,E,M

Hewlett-Packard France
3, Rue Graham Bell
BP 5149
57074 **METZ** Cedex
Tel: (87) 36-13-31
Telex: 860602F
C,E

Hewlett-Packard France
Miniparc-ZIRST
Chemin du Vieux Chêne
38240 **MEYLAN** (Grenoble)
Tel: (76) 90-38-40
980124 HP Grenobe
C

Hewlett-Packard France
Bureau vert du Bois Briand
Cheman de la Garde
- CP 212 212
44085 **NANTES** Cedex
Tel: (40) 50-32-22
Telex: 711085F
A,C,E,CM*,P

Hewlett-Packard France
125, Rue du Faubourg Bannier
45000 **ORLEANS**
Tel: 33-38-62-2031
E,P*

Hewlett-Packard France
Zone Industrielle de Courtaboeuf
Avenue des Tropiques
91947 **LES ULIS** Cedex (Orsay)
Tel: 33-6-907 7825
Telex: 600048F
A,C,CM,E,M,P**

Hewlett-Packard France
15, Avenue de L'Amiral-Bruix
75782 **PARIS** Cedex 16
Tel: 33-15-02-1220
Telex: 613663F
C,P*

Hewlett-Packard France
242 Ter, Ave J Mermoz
64000 **PAU**
Tel: 33-59-80-3802
Telex: 550365F
C,E*

Hewlett-Packard France
6, Place Sainte Croix
86000 **POITIERS**
Tel: 33-49-41-2707
Telex: 792335F
C, E*

Hewlett-Packard France
47, Rue de Chativesle
51100 **REIMS**
Tel: 33-26-88-6919
C, P*

Hewlett-Packard France
Parc d'activités de la Poterie
Rue Louis Kerautret-Botmel
35000 **RENNES**
Tel: 33-99-51-4244
Telex: 740912F
A*,C,E,M,P*

Hewlett-Packard France
98 Avenue de Bretagne
76100 **ROUEN**
Tel: 33-35-63-5766
Telex: 770035F
C,E

Hewlett-Packard France
4, Rue Thomas-Mann
Boite Postale 56
67033 **STRASBOURG** Cedex
Tel: (88) 28-56-46
Telex: 890141F
C,E,M,P*

Hewlett-Packard France
Le Péripole III
3, Chemin du Pigeonnier de la Cépière
31081 **TOULOUSE** Cedex
Tel: 33-61-40-1112
Telex: 531639F
A,C,E,M,P*

Hewlett-Packard France
Les Cardoulines
Batiment B2
Route des Dolines
Parc d'activite de Valbonne
Sophia Antipolis
06560 **VALBONNE** (Nice)
Tel: (93) 65-39-40
C

Hewlett-Packard France
9, Rue Baudin
26000 **VALENCE**
Tel: 33-75-42-7616
C**

Hewlett-Packard France
Carolor
ZAC de Bois Briand
57640 **VIGY** (Metz)
Tel: (8) 771 20 22
C

Hewlett-Packard France
Parc d'activité des Prés
1, Rue Papin Cedex
59658 **VILLENEUVE D'ASCQ**
Tel: 33-20-91-4125
Telex: 160124F
C,E,M,P

Hewlett-Packard France
Parc d'activités Paris-Nord 11
Boite Postale 60020
95971 Roissy Charles de Gaulle
VILLEPINTE
Tel: (1) 48 63 80 80
Telex: 211032F
C,E,M,P*

GABON

Sho Gabon
P.O. Box 89
LIBREVILLE
Tel: 721 484
Telex: 5230

GERMAN FEDERAL REPUBLIC

Hewlett-Packard GmbH
Vertriebszentrum Mitte
Hewlett-Packard-Strasse
D-6380 **BAD HOMBURG**
Tel: (06172) 400-0
Telex: 410 844 hpbhg
A,C,E,M,P

Hewlett-Packard GmbH
Geschäftsstelle
Keithstrasse 2-4
D-1000 **BERLIN** 30
Tel: (030) 21 99 04-0
Telex: 018 3405 hpbln d
A,C,E,M,P

Hewlett-Packard GmbH
Verbindungsstelle Bonn
Friedrich-Ebert-Allee 26
5300 **BONN**
Tel: (0228) 234001
Telex: 8869421

Hewlett-Packard GmbH
Vertriebszentrun Sudwest
Schickardstrasse 2
D-7030 **BOBLINGEN**
Postfach 1427
Tel. (07031) 645-0
Telex: 7265 743 hep
A,C,CM,E,M,P

Hewlett-Packard GmbH
Zentralbereich Mktg
Herrenberger Strasse 130
D-7030 **BOBLINGEN**
Tel: (07031) 14-0
Telex: 7265739 hep

Hewlett-Packard GmbH
Geschäftsstelle
Schleefstr. 28a
D-4600 **DORTMUND**-41
Tel: (0231) 45001
Telex: 822858 hepdod
A,C,E

Hewlett-Packard gmbH
Reparaturzentrum Frankfurt
Berner Strasse 117
6000 **FRANKFURT/MAIN 56**
Tel: (069) 500001-0
Telex: 413249 hpffm

Hewlett-Packard GmbH
Vertriebszentrum Nord
Kapstadtring 5
D-2000 **HAMBURG** 60
Tel: 49-40-63-804-0
Telex: 021 63 032 hphh d
A,C,E,M,P

Hewlett-Packard GmbH
Geschäftsstelle
Heidering 37-39
D-3000 **HANNOVER** 61
Tel: (0511) 5706-0
Telex: 092 3259 hphan
A,C,CM,E,M,P

Hewlett-Packard GmbH
Geschäftsstelle
Rosslauer Weg 2-4
D-6800 **MANNHEIM**
Tel: 49-0621-70-05-0
Telex: 0462105 hpmhm
A,C,E

Hewlett-Packard GmbH
Geschäftsstelle
Messerschmittstrasse 7
D-7910 **NEU ULM**
Tel: 49-0731-70-73-0
Telex: 0712816 HP ULM-D
A,C,E*

Hewlett-Packard GmbH
Geschäftsstelle
Emmericher Strasse 13
D-8500 **NURNBERG** 10
Tel: (0911) 5205-0
Telex: 0623 860 hpnbg
C,CM,E,M,P

Hewlett-Packard GmbH
Vertriebszentrum Ratingen
Berliner Strasse 111
D-4030 **RATINGEN** 4
Postfach 31 12
Tel. (02102) 494-0
Telex: 589 070 hprad
A,C,E,M,P

Hewlett-Packard GmbH
Vertriebszentrum Muchen
Eschenstrasse 5
D-8028 **TAUFKIRCHEN**
Tel. 49-89-61-2070
Telex: 0524985 hpmch
A,C,CM,E,M,P

Hewlett-Packard GmbH
Geschäftsstelle
Ermisallee
7517 **WALDBRONN** 2
Postfach 1251
Tel. (07243) 602-0
Telex: 782 838 hepk
A,C,E

GREAT BRITAIN
See United Kingdom

GREECE

Hewlett-Packard A.E.
178, Kifissias Avenue
6th Floor
Halandri-**ATHENS**
Greece
Tel. 301116473 360, 301116726 090
Telex: 221 286 HPHLGR
A,C,CM**,E,M,P

Kostas Karaynnis S.A.
8, Omirou Street
ATHENS 133
Tel. 32 30 303, 32 37 371
Telex: 215962 RKAR GR
A,C*,CM,E

Impexin
Intelect Div
209 Mesogion
11525 **ATHENS**
Tel. 6474481/2
Telex. 216286
P

Hanri Company
38, Mihalakopoulou
ATHENS 612
Tel. 7236071
Telex: 218767
M*

Heliamco
P.O. Box 87528
18507 **PIRAEUS**
Tel. 4827049
Telex: 241441
A

GUATEMALA

IPESA DE GUATEMALA
Avenida Reforma 3-48, Zona 9
GUATEMALA CITY
Tel. 316627, 317853,66471/5
9-011-502-2-316627
Telex: 3055765 IPESA GU
A,C,CM,E,M,P

HONG KONG

Hewlett-Packard Hong Kong, Ltd
G.P.O. Box 795
5th Floor, Sun Hung Kai Centre
30 Harbour Road, Wan Chai
HONG KONG
Tel: 852-5-832-3211
Telex: 66678 HEWPA HX
Cable: HEWPACK HONG KONG
E,C,P

CET Ltd
10th Floor, Hua Asia Bldg
64-66 Gloucester Road
HONG KONG
Tel: (5) 200922
Telex: 85148 CET HX
CM

Schmidt & Co. (Hong Kong) Ltd
18th Floor, Great Eagle Centre
23 Harbour Road, Wanchai
HONG KONG
Tel: 5-8330222
Telex: 74766 SCHMC HX
A,M

ICELAND

Hewlett-Packard Iceland
Hoefdabakka 9
112 **REYKJAVIK**
Tel: 354-1-67-1000
Telex: 37409
A,C,CM,E,M,P

INDIA

Computer products are sold through Blue Star Ltd. All computer repairs and maintenance service is done through Computer Maintenance Corp.

Blue Star Ltd
B.D. Patel House
Near Sardar Patel Colony
AHMEDABAD 380 014
Tel: 403531, 403532
Telex: 0121-234
Cable: BLUE FROST
A,C,CM,E

Blue Star Ltd
40/4 Lavelle Road
BANGALORE 560 001
Tel: 57881, 867780
Telex: 0845-430 BSLBIN
Cable: BLUESTAR
A,C*,CM,E

Blue Star Ltd.
Band Box House
Prabhadevi
BOMBAY 400 025
Tel: 4933101, 4933222
Telex: 011-71051
Cable: BLUESTAR
A,M

Blue Star Ltd
Sahas
414/2 Vir Savarkar Marg
Prabhadevi
BOMBAY 400 025
Tel: 422-6155
Telex: 011-71193 BSSS IN
Cable: FROSTBLUE
A,CM,E,M

Blue Star Ltd
Kalyan, 19 Vishwas Colony
Alkapuri, **BORODA**, 390 005
Tel: 65235, 65236
Cable: BLUE STAR
A

Blue Star Ltd.
7 Hare Street
P.O. Box 506
CALCUTTA 700 001
Tel: 230131, 230132
Telex: 031-61120 BSNF IN
Cable: BLUESTAR
A,M,C,E

Blue Star Ltd
133 Kodambakkam High Road
MADRAS 600 034
Tel: 472056, 470238
Telex: 041-379
Cable: BLUESTAR
A,M

Blue Star Ltd
13 Community Center
New Friends Colony
NEW DELHI 110 065
Tel: 682547
Telex: 031-2463
Cable: BLUEFROST
A,C*,CM,E,M

Blue Star Ltd.
15/16 C Wellesley Rd.
PUNE 411 011
Tel: 22775
Cable: BLUE STAR
A

Blue Star Ltd
2-2-47/1108 Bolarum Rd.
SECUNDERABAD 500 003
Tel: 72057, 72058
Telex: 0155-459
Cable: BLUEFROST
A,C,E

Blue Star Ltd.
T.C. 7/603 Poornima
Maruthunkuzhi
TRIVANDRUM 695 013
Tel: 65799, 65820
Telex: 0884-259
Cable: BLUESTAR
E

Computer Maintenance Corporation Ltd
115, Sarojini Devi Road
SECUNDERABAD 500 003
Tel: 310-184, 345-774
Telex: 031-2960
C**

INDONESIA

BERCA Indonesia P.T.
P.O.Box 496/Jkt
Jl. Abdul Muis 62
JAKARTA
Tel: 21-373009
Telex: 46748 BERSAL IA
Cable: BERSAL JAKARTA
P

BERCA Indonesia P.T.
P.O. Box 2497/Jkt
Antara Bldg., 12th Floor
Jl. Medan Merdeka Selatan 17
JAKARTA-PUSAT
Tel: 21-340417
Telex: 46748 BERSAL IA
A,C,E,M,P

BERCA Indonesia P.T.
Jalan Kutai 24
SURABAYA
Tel: 67118
Telex: 31146 BERSAL SB
Cable: BERSAL-SURABAYA
A*,E,M,P

IRAQ

Hewlett-Packard Trading S.A.
Service Operation
Al Mansoor City 9B/3/7
BAGHDAD
Tel: 551-49-73
Telex: 212-455 HEPAIRAQ IK
C

IRELAND

Hewlett-Packard Ireland Ltd.
Temple House, Temple Road
Blackrock, Co. **DUBLIN**
Tel: 88/333/99
Telex: 30439
C,E,P

Hewlett-Packard Ltd
75 Belfast Rd, Carrickfergus
Belfast BT38 8PH
NORTHERN IRELAND
Tel: 09603-67333
Telex: 747626
M

ISRAEL

Eidan Electronic Instrument Ltd
P.O. Box 1270
JERUSALEM 91000
16, Ohaliav St.
JERUSALEM 94467
Tel: 533 221, 553 242
Telex: 25231 AB/PAKRD IL
A,M

Computation and Measurement Systems (CMS) Ltd
11 Masad Street
67060
TEL-AVIV
Tel: 388 388
Telex: 33569 Motil IL
C,CM,E,P

ITALY

Hewlett-Packard Italiana S.p.A.
Traversa 99C
Via Giulio Petroni, 19
I-70124 **BARI**
Tel: (080) 41-07-44
C,M

Hewlett-Packard Italiana S.p.A.
Via Emilia, 51/C
I-40011 **BOLOGNA** Anzola Dell'Emilia
Tel: 39-051-731061
Telex: 511630
C,E,M

Hewlett-Packard Italiana S.p.A.
Via Principe Nicola 43G/C
I-95126 **CATANIA**
Tel: (095) 37-10-87
Telex: 970291
C

Hewlett-Packard Italiana S.p.A.
Via G. di Vittorio 10
20094 **CORSICO** (Milano)
Tel: 39-02-4408351

Hewlett-Packard Italiana S.p.A.
Viale Brigata Bisagno 2
16129 **GENOVA**
Tel: 39-10-541141
Telex: 215238

Hewlett-Packard Italiana S.p.A.
Viale G. Modugno 33
I-16156 **GENOVA PEGLI**
Tel: (010) 68-37-07
Telex: 215238
C,E

Hewlett-Packard Italiana S.p.A.
Via G. di Vittorio 9
I-20063 **CERNUSCO SUL NAVIGLIO**
(Milano)
Tel: (02) 923691
Telex: 334632
A,C,CM,E,M,P

Hewlett-Packard Italiana S.p.A.
Via Nuova Rivoltana 95
20090 **LIMITO** (Milano)
Tel: 02-92761

Hewlett-Packard Italiana S.p.A.
Via Nuova San Rocco a Capodimonte, 62/A
I-80131 **NAPOLI**
Tel: (081) 7413544
Telex: 710698
A**,C,E,M

Hewlett-Packard Italiana S.p.A.
Via Orazio 16
80122 **NAPOLI**
Tel: (081) 7611444
Telex: 710698

Hewlett-Packard Italiana S.p.A.
Via Pelizzo 15
35128 **PADOVA**
Tel: 39-49-664-888
Telex: 430315
A,C,E,M

Hewlett-Packard Italiana S.p.A.
Viale C. Pavese 340
I-00144 **ROMA EUR**
Tel: 39-65-48-31
Telex: 610514
A,C,E,M,P*

Hewlett-Packard Italiana S.p.A.
Via di Casalina 57/C
500518 **SCANDICCI-FIRENZE**
Tel: 39-55-753863
C,E,M

Hewlett-Packard Italiana S.p.A.
Corso Svizzera, 185
I-10144 **TORINO**
Tel: 39-11-74-4044
Telex: 221079
A*,C,E

IVORY COAST

S.I.T.E.L.
Societe Ivorienne de Telecommunications
Bd. Giscard d'Estaing
Carrefour Marcory
Zone 4 A.
Borte postale 2580
ABIDJAN 01
Tel: 353600
Telex: 43175
E

S.I.T.I.
Immeuble "Le General"
Av. du General de Gaulle
01 BP 161
ABIDJAN 01
Tel: 321227
Telex: 22149
C,P

JAPAN

Yokogawa-Hewlett-Packard Ltd.
152-1, Onna
ATSUGI, Kanagawa, 243
Tel: (0462) 25-0031
C,CM,E

Yokogawa-Hewlett-Packard Ltd.
Meiji-Seimei Bldg. 6F
3-1 Motochiba-Cho
CHIBA, 280
Tel: (0472) 25 7701
C,E

Yokogawa-Hewlett-Packard Ltd.
Yasuda-Semei Hiroshima Bldg.
6-11, Hon-dori, Naka-ku
HIROSHIMA, 730
Tel: (082) 241-0611

Yokogawa-Hewlett-Packard Ltd.
Towa Building
2-2-3 Kaigan-dori, Chuo-ku
KOBE, 650
Tel: (078) 392-4791
C,E

Yokogawa-Hewlett-Packard Ltd.
Kumagaya Asahi 82 Bldg.
3-4 Tsukuba
KUMAGAYA, Saitama 360
Tel: (0485) 24-6563
C,CM,E

Yokogawa-Hewlett-Packard Ltd
Asahi Shinbun Daiichi Semei Bldg.
4-7, Hanabata-cho
KUMAMOTO, 860
Tel: 96-354-7311
C,E

Yokogawa-Hewlett-Packard Ltd.
Shin-Kyoto Center Bldg.
614, Higashi-Shiokoji-cho
Karasuma-Nishiiru
KYOTO, 600
Tel: 075-343-0921
C,E

Yokogawa-Hewlett-Packard Ltd.
Mito Mitsui Bldg.
1-4-73, Sanno-maru
MITO, Ibaraki 310
Tel: (0292) 25-7470
C,CM,E

Yokogawa-Hewlett-Packard Ltd.
Meiji-Semei Kokubun Bldg.
7-8 Kokubun, 1 Chome, Sendai
MIYAGI, 980
Tel: (0222) 25-1011
C,E

Yokogawa-Hewlett-Packard Ltd.
Gohda Bldg. 2F
1-2-10 Gohda Okaya-Shi
Okaya-Shi
NAGANO, 394
Tel: (0266) 23 0851
C,E

Yokogawa-Hewlett-Packard Ltd.
Nagoya Kokusai Center Building
1-47-1, Nagono, Nakamura-ku
NAGOYA, AICHI 450
Tel: (052) 571-5171
C,CM,E,M

Yokogawa-Hewlett-Packard Ltd
Sa-Kyo-Ren Building
1-2 Dote-cho
OOMIYA-SHI SAITAMA 330
Tel: (0486) 45-8031

SALES & SUPPORT OFFICES
Arranged alphabetically by country

JAPAN (Cont'd)

Yokogawa-Hewlett-Packard Ltd
Chuo Bldg., 5-4-20 Nishi-Nakajima
4-20 Nishinakajima, 5 Chome,
Yodogawa-ku
OSAKA, 532
Tel: (06) 304-6021
Telex: YHPOSA 523-3624
C,CM,E,M,P*

Yokogawa-Hewlett-Packard Ltd.
1-27-15, Yabe
SAGAMIHARA Kanagawa, 229
Tel: 0427 59-1311

Yokogawa-Hewlett-Packard Ltd.
Hamamtsu Motoshiro-Cho Daichi
Seimei Bldg 219-21, Motoshiro-Cho
Hamamatsu-shi
SHIZUOKA, 430
Tel: (0534) 56 1771
C,E

Yokogawa-Hewlett-Packard Ltd
Shinjuku Daiichi Seimei Bldg.
2-7-1, Nishi Shinjuku
Shinjuku-ku,**TOKYO** 163
Tel: 03-348-4611
C,E,M

Yokogawa Hewlett-Packard Ltd.
9-1, Takakura-cho
Hachioji-shi, **TOKYO**, 192
Tel: 81-426-42-1231
C,E

Yokogawa-Hewlett-Packard Ltd.
3-29-21 Takaido-Higashi, 3 Chome
Suginami-ku **TOKYO** 168
Tel: (03) 331-6111
Telex: 232-2024 YHPTOK
C,CM,E,P*

Yokogawa Hokushin Electric
Corporation
Shinjuku-NS Bldg. 10F
4-1 Nishi-Shinjuku 2-Chome
Shinjuku-ku
TOKYO, 163
Tel: (03) 349-1859
Telex: J27584
A

Yokogawa Hokushin Electric Corp
9-32 Nakacho 2 Chome
Musashino-shi
TOKYO, 180
Tel: (0422) 54-1111
Telex: 02822-421 YEW MTK J
A

Yokogawa-Hewlett-Packard Ltd
Meiji-Seimei
Utsunomiya Odori Building
1-5 Odori, 2 Chome
UTSUNOMIYA, Tochigi 320
Tel: (0286) 33-1153
C,E

Yokogawa-Hewlett-Packard Ltd.
Yasuda Seimei Nishiguchi Bldg.
30-4 Tsuruya-cho, 3 Chome
Kanagawa-ku, **YOKOHAMA** 221
Tel: (045) 312-1252
C,CM,E

JORDAN

Scientific and Medical Supplies Co.
P.O. Box 1387
AMMAN
Tel: 24907, 39907
Telex: 21456 SABCO JO
C,E,M,P

KENYA

ADCOM Ltd., Inc., Kenya
P.O.Box 30070
NAIROBI
Tel: 331955
Telex: 22639
E,M

KOREA

Samsung Hewlett-Packard Co. Ltd.
Dongbang Yeoeuido Building
12-16th Floors
36-1 Yeoeuido-Dong
Youngdeungpo-Ku
SEOUL
Tel: 784-4666, 784-2666
Telex: 25166 SAMSAN K
C,CM,E,M,P

Young In Scientific Co., Ltd
Youngwha Building
547 Shinsa Dong, Kangnam-Ku
SEOUL, 135
Tel: 546-7771
Telex: K23457 GINSCO
A

Dongbang Healthcare
Products Co. Ltd.
Suite 301 Madical Supply Center
Bldg. 1-31 Dongsungdong
Jong Ro-gu, **SEOUL**
Tel: 764-1171, 741-1641
Telex: K25706 TKBKO
Cable: TKBEEPKO
M

KUWAIT

Al-Khaldiya Trading & Contracting
P.O. Box 830
SAFAT
Tel: 424910, 411726
Telex: 22481 AREEG KT
Cable: VISCOUNT
E,M,A

Gulf Computing Systems
P.O. Box 25125
SAFAT
Tel: 435969
Telex: 23648
P

Photo & Cine Equipment
P.O. Box 270
SAFAT
Tel: 2445111
Telex: 22247 MATIN KT
Cable: MATIN KUWAIT
P

W.J. Towell Computer Services
P.O. Box 5897
SAFAT
Tel: 2462640/1
Telex: 30336 TOWELL KT
C

LEBANON

Computer Information Systems S.A.L.
Chammas Building
P.O. Box 11-6274 Dora
BEIRUT
Tel: 89 40 73
Telex: 42309 chacis le
C,E,M,P

LIBERIA

Unichemicals Inc
P.O. Box 4509
MONROVIA
Tel: 224282
Telex: 4509
E

LUXEMBOURG

Hewlett-Packard Belgium S.A./N.V.
Blvd de la Woluwe, 100
Woluwedal
B-1200 **BRUSSELS**
Tel: (02) 762-32-00
Telex: 23-494 paloben bru
A,C,CM,E,M,P

MADAGASCAR

Technique et Precision
12, rue de Nice
P.O. Box 1227
101 **ANTANANARIVO**
Tel: 22090
Telex: 22255
P

MALAYSIA

Hewlett-Packard Sales (Malaysia)
Sdn. Bhd.
9th Floor
Chung Khiaw Bank Building
46, Jalan Raja Laut
50736 **KUALA LUMPUR, MALAYSIA**
Tel: 03-2986555
Telex: 31011 HPSM MA
A,C,E,M,P*

Protel Engineering
P.O.Box 1917
Lot 6624, Section 64
23/4 Pending Road
Kuching, **SARAWAK**
Tel: 36299
Telex: 70904 PROMAL MA
Cable: PROTELENG
A,E,M

MALTA

Philip Toledo Ltd.
Kirkirkara P.O. Box 11
Notabile Rd.
MRIEHEL
Tel: 447 47, 455 66, 4915 25
Telex: Media MW 649
E,M,P

MAURITIUS

Blanche Birger Co. Ltd.
18, Jules Koenig Street
PORT LOUIS
Tel: 20828
Telex: 4296
P

MEXICO

Hewlett-Packard de Mexico,
S.A. de C.V.
Rio Nilo No. 4049 Desp. 12
Fracc. Cordoba
JUAREZ
Tel: 161-3-15-62
P

Hewlett-Packard de Mexico,
S.A. de C.V.
Condominio Kadereyta
Circuito del Mezon No. 186 Desp. 6
COL. DEL PRADO - 76030 Qro
Tel: 463-6-02-71
P

Hewlett-Packard de Mexico,
S.A. de C.V.
Monti Morelos No. 299
Fraccionamiento Loma Bonita 45060
GUADALAJARA, Jalisco
Tel: 36-31-48-00
Telex: 0684 186 ECOME
P

Microcomputadoras
Hewlett-Packard, S.A.
Monti Pelvoux 115
LOS LOMAS, Mexico, D.F.
Tel: 520-9127
P

Microcomputadoras Hewlett-Packard,
S.A. de C.V.
Monte Pelvoux No. 115
Lomas de Chapultepec, 11000
MEXICO, D.F.
Tel: 520-9127
P

Hewlett-Packard de Mexico,
S.A. de C.V.
Monte Pelvoux No. 111
Lomas de Chapultepec
11000 **MEXICO, D.F.**
Tel: 5-40-62-28, 72-66, 50-25
Telex: 17-74-507 HEWPACK MEX
A,C,CM,E,M,P

Hewlett-Packard De Mexico (Polanco)
Avenida Ejercito Nacional #579
2da/3er piso
Colonia Granada 11560
MEXICO D.F.
Tel: 254-4433
P

Hewlett-Packard de Mexico,
S.A. de C.V.
Czda. del Valle
409 Ote. 4th Piso
Colonia del Valle
Municipio de Garza
Garcia Nuevo Leon
66220 **MONTERREY**, Nuevo León
Tel: 83-78-42-40
Telex: 382410 HPMY
C

Infograficas y Sistemas
del Noreste, S.A.
Rio Orinoco #171 Oriente
Despacho 2001
Colonia Del Valle
MONTERREY
Tel: 559-4415, 575-3837
Telex: 483164
A,E

Hewlett-Packard de Mexico,
S.A. de C.V.
Blvd. Independencia No. 2000 Ote.
Col. Estrella
TORREON, COAH.
Tel: 171-18-21-99
P

MOROCCO

Etablissement Hubert Dolbeau & Fils
81 rue Karatchi
B.P. 11133
CASABLANCA
Tel: 3041-82, 3068-38
Telex: 23051, 22822
E

Gerep
2, rue Agadir
Boite Postale 156
CASABLANCA 01
Tel: 272093, 272095
Telex: 23 739
P

Sema-Maroc
Dept. Seric
6, rue Lapebie
CASABLANCA
Tel: 260980
Telex: 21641
C,P

NETHERLANDS

Hewlett-Packard Nederland B.V.
Startbaan 16
NL-1187 XR **AMSTELVEEN**
P.O. Box 667
NL-1180 AR **AMSTELVEEN**
Tel: (020) 547-6911
Telex: 13 216 HEPA NL
A,C,CM,E,M,P

Hewlett-Packard Nederland B.V.
Bongerd 2
P.O. Box 41
NL 2900AA **CAPELLE A/D IJSSEL**
Tel: 31-20-51-6444
Telex: 21261 HEPAC NL
C,E

Hewlett-Packard Nederland B.V.
Pastoor Petersstraat 134-136
P.O. Box 2342
NL 5600 CH **EINDHOVEN**
Tel: 31-40-32-6911
Telex: 51484 hepae nl
C,E,P

NEW ZEALAND

Hewlett-Packard (N.Z.) Ltd
5 Owens Road
P.O. Box 26-189
Epsom, **AUCKLAND**
Tel: 64-9-687-159
Cable: HEWPAK Auckland
C,CM,E,P*

Hewlett-Packard (N.Z.) Ltd.
184-190 Willis Street
WELLINGTON
P.O. Box 9443
Courtenay Place, **WELLINGTON 3**
Tel: 64-4-887-199
Cable: HEWPACK Wellington
C,CM,E,P

Northrop Instruments & Systems Ltd
369 Khyber Pass Road
P.O. Box 8602
AUCKLAND
Tel: 794-091
Telex: 60605
A,M

Northrop Instruments & Systems Ltd
110 Mandeville St.
P O Box 8388
CHRISTCHURCH
Tel 488-873
Telex 4203
A,M

Northrop Instruments & Systems Ltd.
Sturdee House
85-87 Ghuznee Street
P O Box 2406
WELLINGTON
Tel 850-091
Telex NZ 3380
A,M

NIGERIA

Elmec Nigeria Ltd.
45 Saka Tinubu St
Victoria Island
LAGOS
Tel 61-98-94
Telex 20-117
E

NORTHERN IRELAND
See United Kingdom

NORWAY

Hewlett-Packard Norge A/S
Folke Bernadottes vei 50
P O Box 3558
N-5033 **FYLLINGSDALEN** (Bergen)
Tel 0047/5/16 55 40
Telex 76621 hpnas n
C,E,M

Hewlett-Packard Norge A/S
Osterndalen 16-18
P O Box 34
N-1345 **OESTERAAS**
Tel 47-2-17-1180
Telex 76621 hpnas n
A,C,CM,E,M,P

Hewlett-Packard Norge A/S
Boehmergt 42
Box 2470
N-5037 **SOLHEIMSVIK**
Tel 0047/5/29 00 90

OMAN

Khimji Ramdas
P O Box 19
MUSCAT/SULTANATE OF OMAN
Tel 795 901
Telex 3489 BROKER MB MUSCAT
P

Suhail & Saud Bahwan
P O Box 169
MUSCAT/SULTANATE OF OMAN
Tel 734 201-3
Telex 5274 BAHWAN MB
E

Imtac LLC
P O Box 9196
MINA AL FAHAL/SULTANATE OF OMAN
Tel 70-77-27, 70-77-23
Telex 3865 Tawoos On
A,C,M

PAKISTAN

Mushko & Company Ltd
House No 16, Street No 16
Sector F-6/3
ISLAMABAD
Tel 824545
Telex 54001 Muski Pk
Cable FEMUS Islamabad
A,E,P*

Mushko & Company Ltd.
Oosman Chambers
Abdullah Haroon Road
KARACHI 0302
Tel: 524131, 524132
Telex: 2894 MUSKO PK
Cable: COOPERATOR Karachi
A,E,P*

PANAMA

Electronico Balboa, S.A.
Calle Samuel Lewis, Ed. Alfa
Apartado 4929
PANAMA CITY
Tel 9-011-507-636613
Telex: 368 3483 ELECTRON PG
CM,E,M,P

PERU

Cia Electro Médica S.A.
Los Flamencos 145, Ofc. 301/2
San Isidro
Casilla 1030
LIMA 1
Tel 9-011-511-4-414325, 41-3705
Telex 39425257 PE PB SIS
CM,E,M,P

SAMS S.A.
Arenida Republica de Panama 3534
San Isidro, **LIMA**
Tel 9-011-511-4-229332/413984/413226
Telex 39420450 PE LIBERTAD
A,C,P

PHILIPPINES

The Online Advanced Systems Corp
2nd Floor, Electra House
115-117 Esteban Street
P O Box 1510
Legaspi Village, Makati
Metro **MANILA**
Tel 815-38-10 (up to 16)
Telex 63274 ONLINE PN
A,C,E,M,P

PORTUGAL

Mundinter Intercambio
Mundial de Comercio S.A.R.L.
Av. Antonio Augusto Aguiar 138
Apartado 2761
LISBON
Tel (19) 53-21-31, 53-21-37
Telex 16691 munter p
M

Soquimica
Av. da Liberdade, 220-2
1298 **LISBOA** Codex
Tel 56-21-82
Telex 13316 SABASA
A

Telectra-Empresa Técnica de Equipmentos Eléctricos S.A.R.L.
Rua Rodrigo da Fonseca 103
P.O. Box 2531
LISBON 1
Tel (19) 68-60-72
Telex 12598
CM,E

C.P.C.S.I
Rua de Costa Cabral 575
4200 **PORTO**
Tel 499174/495173
Telex 26054
C,P

PUERTO RICO

Hewlett-Packard Puerto Rico
101 Munoz Rivera Av.
Esq. Calle Ochoa
HATO REY, Puerto Rico 00918
Tel: (809) 754-7800
A,C,CM,M,E,P

QATAR

Computer Arabia
P O Box 2750
DOHA
Tel: 428555
Telex: 4806 CHPARB
P

Nasser Trading & Contracting
P.O.Box 1563
DOHA
Tel 422170
Telex: 4439 NASSER DH
M

SAUDI ARABIA

Modern Electronics Establishment
Hewlett-Packard Division
P O Box 281
Thuobah
AL-KHOBAR 31952
Tel: 895-1780, 895-1764
Telex: 671 106 HPMEEK SJ
Cable: ELECTA AL-KHOBAR
C,E,M

Modern Electronics Establishment
Hewlett-Packard Division
P O Box 1228
Redec Plaza, 6th Floor
JEDDAH
Tel 644 96 28
Telex: 4027 12 FARNAS SJ
Cable: ELECTA JEDDAH
A,C,CM,E,M,P

Modern Electronics Establishment
Hewlett-Packard Division
P O Box 22015
RIYADH 11495
Tel: 491-97 15, 491-63 87
Telex: 202049 MEERYD SJ
C,E,M

Abdul Ghani El Ajou Corp.
P O Box 78
RIYADH
Tel: 40 41 717
Telex: 200 932 EL AJOU
P

SCOTLAND
See United Kingdom

SENEGAL

Societe Hussein Ayad & Cie.
76, Avenue Georges Pompidou
B.P. 305
DAKAR
Tel: 32339
Cable: AYAD-Dakar
E

Moneger Distribution S.A.
1, Rue Parent
B.P. 148
DAKAR
Tel. 215 671
Telex. 587
P

Systeme Service Conseil (SSC)
14, Avenue du Parachoes
DAKAR ETOILE
Tel 219976
Telex. 577
C,P

SINGAPORE

Hewlett-Packard Singapore (Sales) Pte. Ltd.
1150 Depot Road
SINGAPORE, 0410
Tel: 4731788
Telex: 34209 HPSGSO RS
Cable: HEWPACK, Singapore
A,C,E,M,P

Dynamar International Ltd.
Unit 05-11 Block 6
Kolam Ayer Industrial Estate
SINGAPORE 1334
Tel. 747-6188
Telex: 26283 RS
CM

SOUTH AFRICA

Hewlett-Packard So Africa (Pty.) Ltd.
P.O. Box 120
Howard Place, **CAPE PROVINCE**
7450 South Africa
Tel 27 121153-7954
Telex: 57-20006
A,C,CM,E,M,P

Hewlett-Packard So Africa (Pty.) Ltd
2nd Floor Juniper House
92 Overport Drive
DURBAN 4067
Tel. 27-31-28-4178
Telex: 6-22954
C

Hewlett-Packard So Africa (Pty.) Ltd
Shop 6 Linton Arcade
511 Cape Road
Linton Grange
PORT ELIZABETH 6001
Tel 27141130 1201
Telex. 24-2916
C

Hewlett-Packard So Africa (Pty.) Ltd
Fountain Center
Kalkoen Str
Monument Park Ext 2
PRETORIA 0105
Tel: (012) 45 5725
Telex 32163
C,E

Hewlett-Packard So Africa (Pty.) Ltd.
Private Bag Wendywood
SANDTON 2144
Tel 27-11-802-5111, 27-11-802-5125
Telex: 4-20877 SA
Cable: HEWPACK Johannesburg
A,C,CM,E,M,P

SPAIN

Hewlett-Packard Española, S.A.
Calle Entenza, 321
E.-BARCELONA 29
Tel 3/322 24 51, 321 73 54
Telex: 52603 hpbee
A,C,E,M,P

Hewlett-Packard Española, S.A.
Calle San Vicente S/N
Edificio Alboa II-7B
48001 **BILBAO**
Tel: 4/423 83 06
A,C,E,M

Hewlett-Packard Española, S.A.
Crta. N-VI, Km. 16, 400
Las Rozas
E-MADRID
Tel (1) 637 00 11
Telex 23515 HPE
C,M

Hewlett-Packard Española, S.A.
Avda S. Francisco Javier, S/N
Planta 10, Edificio Sevilla 2
E-SEVILLA 5, SPAIN
Tel. 54/64 44 54
Telex: 72933
A,C,M,P

Hewlett-Packard Española, S.A.
Isabel La Catolica, 8
E-46004 **VALENCIA**
Tel 34-6-361 1354
Telex 63435
C,P

Hewlett-Packard Española, S.A.
Av. de Zugazarte, 8
Las Arenas-Guecho
E-48930 **VIZCAYA**
VIZCAYA
Tel. 34-423-83 06
Telex: 33032

SWEDEN

Hewlett-Packard Sverige AB
Östra Tullgatan 3
S-20011 **MALMÖ**
Box 6132
Tel. 46-40-702-70
Telex (854) 17886 (via Spånga office)
C,P

Hewlett-Packard Sverige AB
Elementvagen 16
S-7022 7 **ÖREBRO**
Tel 49-019-10-4820
Telex (654) 17886 (via Spånga office)
C

Hewlett-Packard Sverige AB
Skalholtsgatan 9, Kista
P.O. Box 19
S-16393 **SPÅNGA**
Tel (08) 750-2000
Telex: (854) 17886
Telefax: (08) 7527781
A,C,CM,E,M,P

Hewlett-Packard Sverige AB
Box 266
Topasgatan 1A
S-42123 **VÄSTRA-FRÖLUNDA** (Gothenburg)
Tel 46-031-89-1000
Telex: (854) 17886 (via Spånga office)
A,C,CM,E,M,P

SUDAN

Mediterranean Engineering & Trading Co. Ltd.
P O Box 1025
KHARTOUM
Tel 41184
Telex: 24052
C,P

SWITZERLAND

Hewlett-Packard (Schweiz) AG
Clarastrasse 12
CH-4058 **BASEL**
Tel: 41-61-33-5920
A,C,E,P

Hewlett-Packard (Schweiz) AG
7, rue du Bois-du-Lan
Case postale 365-1366
CH-1217 **MEYRIN** 1
Tel (0041) 22-83-11-11
Telex 27333 HPAG CH
A,C,CM,E,M,P

SALES & SUPPORT OFFICES
Arranged alphabetically by country

SWITZERLAND (Cont'd)

Hewlett-Packard (Schweiz) AG
Allmend 2
CH-8967 **WIDEN**
Tel: 41-57-31-2111
Telex: 53933 hpag ch
Cable: HPAG CH
A,C,CM,E,M,P

Hewlett-Packard (Schweiz) AG
Schwamendingenstrasse 10
CH-8050 **ZURICH**
Tel: 41-1-315-8181
Telex: 823 537 HPAG CH
C,P

SYRIA

General Electronic Inc
Nuri Basha Ahnaf Ebn Kays Street
P.O. Box 5781
DAMASCUS
Tel: 33-24-87
Telex: 44-19-88
Cable: ELECTROBOR DAMASCUS
E

Middle East Electronics
P.O.Box 2308
Abu Rumaneh
DAMASCUS
Tel: 33 45 92
Telex: 411 771 Meesy
M

TAIWAN

Hewlett-Packard Taiwan Ltd.
THM Office
2, Huan Nan Road
CHUNG LI, Taoyuan
Tel: (034) 929-666
C

Hewlett-Packard Taiwan Ltd.
Kaohsiung Office
11/F, 456, Chung Hsiao 1st Road
KAOHSIUNG
Tel: (07) 2412318
C,E

Hewlett-Packard Taiwan Ltd.
8th Floor, Hewlett-Packard Building
337 Fu Hsing North Road
TAIPEI
Tel: (02) 712-0404
Telex: 24439 HEWPACK
Cable: HEWPACK Taipei
A,C,CM,E,M,P

Ing Lih Trading Co
3rd Floor, No. 7, Sect. 2
Jen Ai Road
TAIPEI 100
Tel: (02) 394-8191
Telex: 22894 SANKWANG
A

THAILAND

Unimesa Co. Ltd
30 Patpong Ave., Surwong
BANGKOK 5,
Tel: 235-5727, 234-0991/3
Telex: 84439 Simonco TH
Cable: UNIMESA Bangkok
A,C,E,M

Bangkok Business Equipment Ltd
5/2-6 Dejo Road
BANGKOK
Tel: 234-8670, 234-8671
Telex: 87699-BEQUIPT TH
Cable: BUSIQUIPT Bangkok
P

TOGO

Societe Africaine De Promotion
Immeuble Sagab
Rue d'Atakpame
P.O. Box 4150
LOME
Tel: 21-62-88
Telex: 5357
P

TRINIDAD & TOBAGO

Caribbean Telecoms Ltd.
Corner McAllister Street &
Eastern Main Road, Laventville
P.O. Box 732
PORT-OF-SPAIN
Tel: 624-4213
Telex: 22561 CARTEL WG
Cable: CARTEL, PORT OF SPAIN
CM,E,M,P

Computer and Controls Ltd
P.O. Box 51
1 Taylor Street
PORT-OF-SPAIN
Tel: (809) 622-7719/622-7985
Telex: 38722798 COMCON WG
LOOGO AGENCY 1264
A,P

Feral Assoc.
8 Fitzgerald Lane
PORT-OF-SPAIN
Tel: 62-36864, 62-39255
Telex: 22432 FERALCO
Cable: FERALCO
M

TUNISIA

Tunisie Electronique S.A.R.L.
31 Avenue de la Liberte
TUNIS
Tel: 280-144
C,E,P

Tunisie Electronique S.A.R.L.
94, Av. Jugurtha, Mutuelleville
1002 **TUNIS-BELVEDERE**
Tel: 280144
Telex: 13238
C,E,P

Corema S.A.
1 ter. Av. de Carthage
TUNIS
Tel: 253-821
Telex: 12319 CABAM TN
M

TURKEY

E.M.A.
Mediha Eldern Sokak No. 41/6
Yenisehir
ANKARA
Tel: 319175
Telex: 42321 KTX TR
Cable: EMATRADE ANKARA
M

Teknim Company Ltd.
Iran Caddesi No. 7
Karaklidere
ANKARA
Tel: 275800
Telex: 42155 TKNM TR
C,E

Kurt & Kurt A.S.
Mithatpasa Caddesi No. 75
Kat 4 Kizilay
ANKARA
Tel: 318875/6/7/8
Telex: 42490 MESR TR
A

Saniva Bilgisayar Sistemleri A.S.
Buyukdere Caddesi 103/6
Gayrettepe
ISTANBUL
Tel: 1673180
Telex: 26345 SANI TR
C,P

Best Inc.
Esentepe, Gazeteciler Sitesi
Keskin Kalem
Sokak 6/3, Gayrettepe
ISTANBUL
Tel: 172 1328, 173 3344
Telex: 42490
A

UNITED ARAB EMIRATES

Emitac Ltd.
P.O. Box 1641
SHARJAH
Tel: 591181
Telex: 68136 EMITAC EM
Cable: EMITAC SHARJAH
E,C,M,P,A

Emitac Ltd
P.O. Box 2711
ABU DHABI
Tel: 820419-20
Cable: EMITACH ABUDHABI

Emitac Ltd
P.O. Box 8391
DUBAI,
Tel: 377591

Emitac Ltd
P.O. Box 473
RAS AL KHAIMAH
Tel: 28133, 21270

UNITED KINGDOM
ENGLAND

Hewlett-Packard Ltd
Miller House
The Ring, **BRACKNELL**
Berks RG12 1XN
Tel: 44/344/424-898
Telex: 848733
E

Hewlett-Packard Ltd.
Elstree House, Elstree Way
BOREHAMWOOD, Herts WD6 1SG
Tel: 01 207 5000
Telex: 8952716
C,E

Hewlett-Packard Ltd
Oakfield House, Oakfield Grove
Clifton **BRISTOL**, Avon BS8 2BN
Tel: 44-272-736 806
Telex: 444302
C,E,P

Hewlett-Packard Ltd
9 Bridewell Place
LONDON EC4V 6BS
Tel: 44-01-583-6565
Telex: 298163
C,P

Hewlett-Packard Ltd.
Pontefract Road
NORMANTON, West Yorkshire WF6 1RN
Tel: 44/924/895 566
Telex: 557355
C,P

Hewlett-Packard Ltd.
The Quadrangle
106-118 Station Road
REDHILL, Surrey RH1 1PS
Tel: 44-737-686-55
Telex: 947234
C,E,P

Hewlett-Packard Ltd.
Avon House
435 Stratford Road
Shirley, **SOLIHULL**, West Midlands
B90 4BL
Tel: 44-21-745-8800
Telex: 339105
C,E,P

Hewlett-Packard Ltd.
Heathside Park Road
Cheadle Heath, Stockport
SK3 ORB, United Kingdom
Tel: 44-061-428-0828
Telex: 668068
A,C,E,M,P

Hewlett-Packard Ltd.
Harmon House
No. 1 George Street
UXBRIDGE, Middlesex UX8 1YH
Tel: 895 720 20
Telex: 893134/5
C,CM,E,M,P

Hewlett-Packard Ltd.
King Street Lane
Winnersh, **WOKINGHAM**
Berkshire RG11 5AR
Tel: 44/734/784774
Telex: 8471789
A,C,E,M,P

NORTHERN IRELAND

Hewlett-Packard (Ireland) Ltd.
Carrickfergus Industrial Centre
75 Belfast Road, Carrickfergus
CO. ANTRIM BT38 8PM
Tel: 09603 67333
C,E

Cardiac Services Company
95A Finaghy Road South
BELFAST, BT10 OBY
Tel: 0232-625566
Telex: 747626
M

SCOTLAND

Hewlett-Packard Ltd.
1/3 Springburn Place
College Milton North
EAST KILBRIDE, G74 5NU
Tel: 041-332-6232
Telex: 779615
C,E

Hewlett-Packard Ltd.
SOUTH QUEENSFERRY
West Lothian, EH30 9TG
Tel: 031 331 1188
Telex: 72682 HPSQFYG
C,CM,E,M,P

UNITED STATES

Hewlett-Packard Co
Customer Information Center
Tel: (800) 752-0900
Hours: 6:00 AM to 5:00 PM
Pacific Time

Alabama

Hewlett-Packard Co
2100 Riverchase Center
Building 100 - Suite 118
BIRMINGHAM, AL 35244
Tel: (205) 988-0547
A,C,M,P*

Hewlett-Packard Co
420 Wynn Drive
HUNTSVILLE, AL 35805
Tel: (205) 830-2000
C,CM,E,M*

Alaska

Hewlett-Packard Co
4000 Old Seward Highway
Suite 101
ANCHORAGE, AK 99503
Tel: (907) 563-8855
C,E

Arizona

Hewlett-Packard Co
8080 Pointe Parkway West
PHOENIX, AZ 85044
Tel: (602) 273-8000
A,C,CM,E,M,P

Hewlett-Packard Co
3400 East Britannia Dr
Bldg. C, Suite 124
TUCSON, AZ 85706
Tel: (602) 573-7400
C,E,M**

California

Hewlett-Packard Co
99 South Hill Dr
BRISBANE, CA 94005
Tel: (415) 330-2500
C

Hewlett-Packard Co
1907 North Gateway Blvd
FRESNO, CA 93727
Tel: (209) 252-9652
C,M

Hewlett-Packard Co
1421 S. Manhattan Av.
FULLERTON, CA 92631
Tel: (714) 999-6700
C,CM,E,M

Hewlett-Packard Co
7408 Hollister Ave. #A
GOLETA, CA 93117
Tel: (805) 685-6100
C,E

Hewlett-Packard Co
2525 Grand Avenue
LONG BEACH, CA 90815
Tel: (213) 498-1111
C

Hewlett-Packard Co
5651 West Manchester Ave
LOS ANGELES, CA 90045
Tel: (213) 337-8000

Hewlett-Packard Co
3155 Porter Drive
PALO ALTO, CA 94304
Tel: (415) 857-8000
C,E

Hewlett-Packard Co.
5725 W. Las Positas Blvd
PLEASANTON, CA 94566
Tel (415) 460-0282
C

Hewlett-Packard Co.
4244 So. Market Court, Suite A
SACRAMENTO, CA 95834
Tel (916) 929-7222
A*,C,E,M

Hewlett-Packard Co.
9606 Aero Drive
SAN DIEGO, CA 92123
Tel (619) 279-3200
C,CM,E,M

Hewlett-Packard Co
3003 Scott Boulevard
SANTA CLARA, CA 95054
Tel (408) 988-7000
Telex 910-338-0586
A,C,CM,E

Hewlett-Packard Co
2150 W. Hillcrest Dr
THOUSAND OAKS, CA 91320
(805) 373-7000
C,CM,E

Colorado

Hewlett-Packard Co
2945 Center Green Court South
Suite A
BOULDER, CO 80301
Tel (303) 499-6655
A,C,E

Hewlett-Packard Co
24 Inverness Place, East
ENGLEWOOD, CO 80112
Tel (303) 649-5000
A,C,CM,E,M

Connecticut

Hewlett-Packard Co
500 Sylvan Av.
BRIDGEPORT, CT 06606
Tel (203) 371-6454
C,E

Hewlett-Packard Co
47 Barnes Industrial Road South
WALLINGFORD, CT 06492
Tel (203) 265-7801
A,C,CM,E,M

Florida

Hewlett-Packard Co
2901 N.W. 62nd Street
FORT LAUDERDALE, FL 33309
Tel (305) 973-2600
C,E,M,P*

Hewlett-Packard Co
6800 South Point Parkway
Suite 301
JACKSONVILLE, FL 32216
Tel (904) 636-9955
C*,M**

Hewlett-Packard Co
255 East Drive, Suite B
MELBOURNE, FL 32901
Tel (305) 729-0704
CM,E

Hewlett-Packard Co
6177 Lake Ellenor Drive
ORLANDO, FL 32809
Tel (305) 859-2900
A,C,CM,E,P*

Hewlett-Packard Co
4700 Bayou Blvd
Building 5
PENSACOLA, FL 32503
Tel (904) 476-8422
A,C,M

Hewlett-Packard Co
5550 W. Idlewild, #150
TAMPA, FL 33614
Tel (813) 884-3282
C,E,M,P

Georgia

Hewlett-Packard Co
2015 South Park Place
ATLANTA, GA 30339
Tel (404) 955-1500
Telex 810-766-4890
A,C,CM,E,M,P*

Hewlett-Packard Co
3607 Parkway Lane
Suite 300
NORCROSS, GA 30092
Tel (404) 448-1894
C,E,P

Hawaii

Hewlett-Packard Co.
Pacific Tower
1001 Bishop St
Suite 2400
HONOLULU, HI 96813
Tel (808) 526-1555
A,C,E,M

Idaho

Hewlett-Packard Co
11309 Chinden Blvd
BOISE, ID 83714
Tel (208) 323-2700
C

Illinois

Hewlett-Packard Co.
2205 E. Empire St
P.O. Box 1607
BLOOMINGTON, IL 61702-1607
Tel (309) 662-9411
A,C,E,M**

Hewlett-Packard Co
525 W. Monroe, #1308
CHICAGO, IL 60606
Tel (312) 930-0010
C

Hewlett-Packard Co
1200 East Diehl Road
NAPERVILLE, IL 60566
Tel (312) 357-8800
C

Hewlett-Packard Co.
5201 Tollview Drive
ROLLING MEADOWS, IL 60008
Tel (312) 255-9800
Telex 910-687-1066
A,C,CM,E,M

Indiana

Hewlett-Packard Co
11911 N. Meridian St
CARMEL, IN 46032
Tel (317) 844-4100
A,C,CM,E,M

Hewlett-Packard Co
111 E. Ludwig Road
Suite 108
FT. WAYNE, IN 46825
Tel (219) 482-4283
C,E

Iowa

Hewlett-Packard Co
4070 22nd Av. SW
CEDAR RAPIDS, IA 52404
Tel (319) 390-4250
C,E,M

Hewlett-Packard Co
4201 Corporate Dr
WEST DES MOINES, IA 50265
Tel (515) 224-1435
A**,C,M**

Kansas

Hewlett-Packard Co
North Rock Business Park
3450 N. Rock Rd.
Suite 300
WICHITA, KS 67226
Tel (316) 684-8491
C,E

Kentucky

Hewlett-Packard Co
305 N. Hurstbourne Lane,
Suite 100
LOUISVILLE, KY 40223
Tel (502) 426-0100
A,C,M

Louisiana

Hewlett-Packard Co
160 James Drive East
ST. ROSE, LA 70087
P.O. Box 1449
KENNER, LA 70063
Tel (504) 467-4100
A,C,E,M,P

Maryland

Hewlett-Packard Co
3701 Koppers Street
BALTIMORE, MD 21227
Tel (301) 644-5800
Telex 710-862-1943
A,C,CM,E,M

Hewlett-Packard Co
2 Choke Cherry Road
ROCKVILLE, MD 20850
Tel (301) 948-6370
A,C,CM,E,M

Massachusetts

Hewlett-Packard Co
1775 Minuteman Road
ANDOVER, MA 01810
Tel (617) 682-1500
A,C,CM,E,M,P*

Hewlett-Packard Co
29 Burlington Mall Rd
BURLINGTON, MA 01803-4514
Tel (617) 270-7000
C,E

Michigan

Hewlett-Packard Co
4326 Cascade Road S.E.
GRAND RAPIDS, MI 49506
Tel (616) 957-1970
C,M

Hewlett-Packard Co
39550 Orchard Hill Place Drive
NOVI, MI 48050
Tel (313) 349-9200
A,C,E,M

Hewlett-Packard Co.
560 Kirts Rd
Suite 101
TROY, MI 48084
Tel. (313) 362-5180
C

Minnesota

Hewlett-Packard Co
2025 W. Larpenteur Ave
ST. PAUL, MN 55113
Tel (612) 644-1100
A,C,CM,E,M

Missouri

Hewlett-Packard Co
1001 E. 101st Terrace Suite 120
KANSAS CITY, MO 64131-3368
Tel (816) 941-0411
A,C,CM,E,M

Hewlett-Packard Co
13001 Hollenberg Drive
BRIDGETON, MO 63044
Tel (314) 344-5100
A,C,E,M

Nebraska

Hewlett-Packard
11626 Nicholas St.
OMAHA, NE 68154
Tel (402) 493-0300
C,E,M

New Jersey

Hewlett-Packard Co
120 W. Century Road
PARAMUS, NJ 07652
Tel (201) 265-5000
A,C,CM,E,M

Hewlett-Packard Co
20 New England Av. West
PISCATAWAY, NJ 08854
Tel (201) 562-6100
A,C,CM,E

New Mexico

Hewlett-Packard Co
7801 Jefferson N.E.
ALBUQUERQUE, NM 87109
Tel (505) 823-6100
C,E,M

Hewlett-Packard Co
1362-C Trinity Dr
LOS ALAMOS, NM 87544
Tel (505) 662-6700
C,E

New York

Hewlett-Packard Co
5 Computer Drive South
ALBANY, NY 12205
Tel (518) 458-1550
A,C,E,M

Hewlett-Packard Co
9600 Main Street
CLARENCE, NY 14031
Tel (716) 759-8621
C,E,M

Hewlett-Packard Co
200 Cross Keys Office Park
FAIRPORT, NY 14450
Tel. (716) 223-9950
A,C,CM,E,M

Hewlett-Packard Co
7641 Henry Clay Blvd
LIVERPOOL, NY 13088
Tel (315) 451-1820
A,C,CM,E,M

Hewlett-Packard Co
No. 1 Pennsylvania Plaza
55th Floor
34th Street & 7th Avenue
MANHATTAN NY 10119
Tel (212) 971-0800
C,M*

Hewlett-Packard Co
15 Myers Corner Rd.
Hollowbrook Park, Suite 2D
WAPPINGERS FALLS, NY 12590
Tel (914) 298-9125
CM,E

Hewlett-Packard Co.
2975 Westchester Ave
PURCHASE, NY 10577
Tel (914) 935-6300
C,CM,E

Hewlett-Packard Co.
3 Crossways Park West
WOODBURY, NY 11797
Tel (516) 682-7800
A,C,CM,E,M

North Carolina

Hewlett-Packard Co
305 Gregson Dr.
CARY, NC 27511
Tel (919) 467-6600
C,CM,E,M,P*

Hewlett-Packard Co.
9401 Arrow Point Blvd
Suite 100
CHARLOTTE, NC 28217
Tel (704) 527-8780
C*

Hewlett-Packard Co.
5605 Roanne Way
GREENSBORO, NC 27420
Tel (919) 852-1800
A,C,CM,E,M,P*

Ohio

Hewlett-Packard Co
2717 S. Arlington Road
AKRON, OH 44312
Tel (216) 644-2270
C,E

Hewlett-Packard Co.
4501 Erskine Road
CINCINNATI, OH 45242
Tel (513) 891-9870
C,M

Hewlett-Packard Co.
15885 Sprague Road
CLEVELAND, OH 44136
Tel (216) 243-7300
A,C,CM,E,M

Hewlett-Packard Co.
9080 Springboro Pike
MIAMISBURG, OH 45342
Tel (513) 433-2223
A,C,CM,E*,M

Hewlett-Packard Co.
One Maritime Plaza, 5th Floor
720 Water Street
TOLEDO, OH 43604
Tel (419) 242-2200
C

Hewlett-Packard Co
675 Brooksedge Blvd.
WESTERVILLE, OH 43081
Tel. (614) 891-3344
C,CM,E*

Oklahoma

Hewlett-Packard Co.
3525 N.W. 56th St
Suite C-100
OKLAHOMA CITY, OK 73112
Tel (405) 946-9499
C,E*,M

SALES & SUPPORT OFFICES

Arranged alphabetically by country

UNITED STATES (Cont'd)

Hewlett-Packard Co
6655 South Lewis,
Suite 105
TULSA, OK 74136
Tel: (918) 481-6700
A**,C,E,M*,P*

Oregon

Hewlett-Packard Co
9255 S. W. Pioneer Court
WILSONVILLE, OR 97070
Tel: (503) 682-8000
A,C,E*,M

Pennsylvania

Hewlett-Packard Co
Heatherwood Industrial Park
50 Dorchester Rd
Route 22
HARRISBURG, PA 17112-2799
Tel: (717) 657-5900
C

Hewlett-Packard Co
111 Zeta Drive
PITTSBURGH, PA 15238
Tel: (412) 782-0400
A,C,E,M

Hewlett-Packard Co
2750 Monroe Boulevard
VALLEY FORGE, PA 19482
Tel: (215) 666-9000
A,C,CM,E,M

South Carolina

Hewlett-Packard Co
Brookside Park, Suite 122
1 Harbison Way
COLUMBIA, SC 29212
Tel: (803) 732-0400
C,M

Hewlett-Packard Co
545 N. Pleasantburg Dr
Suite 100
GREENVILLE, SC 29607
Tel: (803) 232-8002
C

Tennessee

Hewlett-Packard Co
One Energy Centr Suite 200
Pellissippi Pkwy
KNOXVILLE, TN 37932
Tel: (615) 966-4747
A,C,E,M,P

Hewlett-Packard Co
3070 Directors Row
Directors Square
MEMPHIS, TN 38131
Tel: (901) 346-8370
A,C,E,M

Hewlett-Packard Co
44 Vantage Way,
Suite 150
NASHVILLE, TN 37228
Tel: (615) 255-1271
A,C,E,M,P

Texas

Hewlett-Packard Co.
1826-P Kramer Lane
AUSTIN, TX 78758
Tel: (512) 835-6771
C,E,P*

Hewlett-Packard Co
5700 Cromo Dr
EL PASO, TX 79912
Tel: (915) 833-4400
C,E*,M**

Hewlett-Packard Co.
3952 Sandshell Drive
FORT WORTH, TX 76137
Tel: (817) 232-9500
C

Hewlett-Packard Co.
10535 Harwin Drive
HOUSTON, TX 77036
Tel: (713) 776-6400
A,C,E,M,P*

Hewlett-Packard Co
3301 West Royal Lane
IRVING, TX 75063
Tel: (214) 869-3377
C,E

Hewlett-Packard Co.
109 E. Toronto, Suite 100
McALLEN, TX 78501
Tel: (512) 630-3030
C

Hewlett-Packard Co.
930 E. Campbell Rd.
RICHARDSON, TX 75081
Tel: (214) 231-6101
A,C,CM,E,M,P*

Hewlett-Packard Co.
1020 Central Parkway South
SAN ANTONIO, TX 78232
Tel: (512) 494-9336
A,C,E,M,P*

Utah

Hewlett-Packard Co
3530 W. 2100 South St
SALT LAKE CITY, UT 84119
Tel: (801) 974-1700
A,C,E,M

Virginia

Hewlett-Packard Co.
840 Greenbrier Circle
Suite 101
CHESAPEAKE, VA 23320
Tel: (804) 424-7105
C,E,M

Hewlett-Packard Co
4305 Cox Road
GLEN ALLEN, VA 23060
Tel: (804) 747-7750
A,C,E,M,P*

Hewlett-Packard Co.
Tanglewood West Bldg.
Suite 240
3959 Electric Road
ROANOKE, VA 24018
Tel: (703) 774-3444
C,E,P

Washington

Hewlett-Packard Co.
15815 S.E. 37th Street
BELLEVUE, WA 98006
Tel: (206) 643-4004
A,C,CM,E,M

Hewlett-Packard Co.
1225 Argonne Rd
SPOKANE, WA 99212
Tel: (509) 922-7000
C

West Virginia

Hewlett-Packard Co.
501 56th Street
CHARLESTON, WV 25304
Tel: (304) 925-0492
A,C,M

Wisconsin

Hewlett-Packard Co.
275 N. Corporate Dr.
BROOKFIELD, WI 53005
Tel: (414) 784-8800
A,C,E*,M

URUGUAY

Pablo Ferrando S.A.C. e I.
Avenida Italia 2877
Casilla de Correo 370
MONTEVIDEO
Tel: 59-82-802-586
Telex: 398802586
A,CM,E,M

Olympia da Uruguay S.A.
Maquinas de Oficina
Avda. del Libertador 1997
Casilla de Correos 6644
MONTEVIDEO
Tel: 91-1809, 98-3807
Telex: 6342 OROU UY
P

VENEZUELA

Hewlett-Packard de Venezuela C.A.
3A Transversal Los Ruices Norte
Edificio Segre 2 & 3
Apartado 50933
CARACAS 1050
Tel: (582) 239-4133
Telex: 251046 HEWPACK
A,C,CM,E,M,P

Hewlett-Packard de Venezuela, C.A.
Centro Ciudad Comercial Tamanaco
Nivel C-2 (Nueva Etapa)
Local 53H05
Chuao, **CARACAS**
Tel: 928291
P

Albis Venezolana S.R.L.
Av. Las Mariles, Ota. Alix,
El Pedregal
Apartado 81025
CARACAS 1080A
Tel: 747984, 742146
Telex: 24009 ALBIS VC
A

Tecnologica Medica del Caribe, C.A.
Multicentro Empresarial del Este
Ave. Libertador
Edif. Libertador
Nucleo "C" - Oficina 51-52
CARACAS
Tel: 339867/333780
M

Hewlett-Packard de Venezuela C.A.
Residencias Tia Betty Local 1
Avenida 3 y con Calle 75
MARACAIBO, Estado Zulia
Apartado 2646
Tel: 58-2-617-5669
Telex: 62464 HPMAR
C,E*

Hewlett-Packard de Venezuela C.A.
Urb. Lomas de Este
Torre Trebol — Piso 11
VALENCIA, Estado Carabobo
Apartado 3347
Tel: (5841) 222992
C,P

YUGOSLAVIA

Do Hermes
General Zdanova 4
YU-11000 **BEOGRAD**
Tel: (011) 342 641
Telex: 11433
A,C,E,M,P

Do Hermes
Celovska 73
YU-61000 **LJUBLJANA**
Tel: (061) 553 170
Telex: 31583
A,C,E,M,P

Elektrotehna
Trdova 51
YU-61000 **LJUBLJANA**
CM

Do Hermes
Kralja Tomislava 1
YU-71000 **SARAJEVO**
Tel: (071) 35 859
Telex: 41634
C**,P

ZAIRE

Computer & Industrial Engineering
25, Avenue de la Justice
B.P. 12797
KINSHASA, Gombe
Tel: 32063
Telex: 21552
C,P

ZAMBIA

R.J. Tilbury (Zambia) Ltd.
P.O. Box 32792
LUSAKA
Tel: 215590
Telex: 40128
E

ZIMBABWE

Field Technical Sales (Private) Limited
45, Kelvin Road North
P.O. Box 3458
SALISBURY
Tel: 705 231
Telex: 4-122 RH
E,P

September 1987

NOTES

NOTES

NOTES

NOTES

NOTES

NOTES

NOTES

Printed in USA 02/91

Customer Order No.
B1013-90009

Manufacturing No.
B1013-91009
Mfg. number is for HP internal use only

B1013-91009