

HP 9000 Series 300 and 800 Computers
**Installing and Administering
NFS Services**



Manual Part Number: B1013-90001
Printed in U.S.A., September 1989

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.

© Copyright 1987, 1988, 1989 Hewlett-Packard Company.

This document contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of Hewlett-Packard Company. The information contained in this document is subject to change without notice.

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).

© Copyright 1980, 1984, AT&T, Inc.

© Copyright 1979, 1980, 1983, The Regents of the University of California.

© Copyright, 1979, 1987, 1988, Sun Microsystems, Inc.

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.

MS-DOS[®] is a U.S. registered trademark of Microsoft Corporation.

UNIX[®] is a U.S. registered trademark of AT&T in the U.S.A. and in other countries.

NFS is a trademark of Sun Microsystems, Inc.

Hewlett-Packard Co.
3404 E. Harmony Rd.
Fort Collins, CO 80525 U.S.A.

Printing History

First Edition September 1989

Contents

Chapter 1: Documentation Overview

| | |
|---------------------------|-----|
| Contents | 1-2 |
| Conventions | 1-5 |
| Documentation Guide | 1-6 |

Chapter 2: NFS Services Overview

| | |
|--|------|
| NFS Services | 2-3 |
| NFS Remote File Access | 2-4 |
| Named Pipes | 2-5 |
| Device Files | 2-6 |
| Remote Execution Facility (REX) | 2-8 |
| Remote Procedure Call (RPC) | 2-8 |
| Remote Procedure Call Protocol Compiler (RPCGEN) | 2-9 |
| External Data Representation (XDR) | 2-10 |
| Network Lock Manager | 2-11 |
| Yellow Pages (YP) | 2-11 |
| YP Advantages | 2-12 |
| YP Disadvantages | 2-12 |
| YP Concepts | 2-13 |
| YP Maps | 2-14 |
| YP Servers and YP Clients | 2-14 |
| YP Domains | 2-15 |
| YP Masters and YP Slaves | 2-16 |
| Virtual Home Environment (VHE) | 2-17 |
| VHE Advantages | 2-17 |
| VHE Disadvantages | 2-19 |
| How VHE Works | 2-20 |
| Example Grouping | 2-21 |

Chapter 3: Installation

| | |
|-------------------------------|------|
| NFS Installation Checklist | 3-1 |
| Key Terms | 3-3 |
| Prepare the HP 9000 System | 3-5 |
| Install the NFS Software | 3-6 |
| Use Update | 3-7 |
| Configure a New Kernel | 3-9 |
| Add a Computer to the Network | 3-10 |
| Relink Programs | 3-11 |

Chapter 4: NFS Configuration and Maintenance

| | |
|---|------|
| Key Terms | 4-2 |
| Guidelines | 4-5 |
| Network Memory | 4-5 |
| Configuration Files | 4-6 |
| Daemons | 4-8 |
| Servers | 4-10 |
| NFS Configuration | 4-11 |
| 1. Compare /etc/newconfig Files to Existing Files | 4-13 |
| 2. Set UIDs and GIDs | 4-14 |
| 3. Create an NFS Server | 4-15 |
| A. Edit /etc/netnfsrc | 4-16 |
| B. Edit /etc/inetd.conf | 4-18 |
| C. Edit /usr/adm/inetd.sec (if necessary) | 4-20 |
| D. Edit /etc/hosts | 4-22 |
| E. Edit /etc/netgroup | 4-23 |
| F. Create and Edit /etc/exports | 4-25 |
| G. Reboot the System (if necessary) | 4-30 |
| 4. Create an NFS Client | 4-31 |
| a. Edit /etc/netnfsrc | 4-31 |
| b. Mount File Systems | 4-34 |
| Mount Guidelines | 4-35 |
| Edit /etc/checklist for Automatic Mounts | 4-40 |
| Execute mount(1M) for Manual Mounts | 4-42 |
| c. Reboot the System (if necessary) | 4-44 |
| 5. Configure YP (optional) | 4-44 |
| 6. Configure VHE (optional) | 4-44 |

**Chapter 4: NFS Configuration and Maintenance
(Continued...)**

| | |
|---|------|
| 7. Execute /etc/netnfsrc | 4-44 |
| Maintenance | 4-45 |
| Remove NFS File Access | 4-45 |
| Unmount File Systems from Client | 4-45 |
| Prevent Access to Server File Systems | 4-47 |
| Update Software | 4-49 |
| Clock Skew | 4-50 |

Chapter 5: Remote Execution Facility (REX)

| | |
|--------------------------------------|------|
| Introduction | 5-1 |
| The on Command | 5-2 |
| Configuration Requirements | 5-4 |
| Environment Simulation | 5-5 |
| Configuring rexd | 5-7 |
| Security Considerations | 5-9 |
| Diagnostics | 5-10 |
| rexid Error Messages | 5-11 |

Chapter 6: The Network Lock Manager

| | |
|--------------------------------------|-----|
| Introduction | 6-1 |
| The Locking Protocol | 6-4 |
| The Network Status Monitor | 6-5 |

Chapter 7: YP Configuration and Maintenance

| | |
|--|------|
| Key Terms | 7-2 |
| YP Databases | 7-5 |
| Local and Global Maps | 7-5 |
| Escape Sequences | 7-7 |
| Netgroups | 7-8 |
| Files Related to YP | 7-11 |
| YP Commands | 7-12 |
| YP Configuration | 7-14 |
| Before Configuring YP | 7-14 |
| Compare /etc/newconfig Files to Existing Files | 7-14 |

Chapter 7: YP Configuration and Maintenance (Continued...)

| | |
|-------------------------------------|------|
| 1. Create a YP Master Server | 7-15 |
| Before Creating a YP Master Server | 7-16 |
| Security | 7-16 |
| Creating a YP Master Server | 7-17 |
| Starting the YP Master Server | 7-18 |
| 2. Create a YP Client | 7-20 |
| Before Creating a YP Client | 7-20 |
| Creating a YP Client | 7-20 |
| Altering a Client's Files | 7-21 |
| Starting the YP Client | 7-25 |
| 3. Create a YP Slave Server | 7-26 |
| Before Creating a YP Slave Server | 7-26 |
| Creating a YP Slave Server | 7-27 |
| Starting the YP Slave Server | 7-28 |
| 4. Propagate YP Maps | 7-29 |
| Verify YP | 7-32 |
| YP Maintenance | 7-33 |
| Disable YP | 7-33 |
| Modify YP Maps | 7-34 |
| Manual Modifications to YP Maps | 7-35 |
| YP Maintenance | 7-36 |
| Add New YP Servers | 7-39 |
| Add New Users to a Node | 7-39 |
| Make a Different Node the YP Master | 7-40 |
| Create or Change YP Password | 7-41 |
| YP Password Guidelines | 7-41 |
| YP Password | 7-42 |
| Log Files | 7-43 |
| Create Non-standard YP Maps | 7-44 |

Chapter 8: VHE Configuration and Maintenance

| | |
|---|-----|
| Configuration Overview | 8-2 |
| Configuration | 8-3 |
| 1. Complete Preparation Steps | 8-3 |
| 2. Compare /etc/newconfig Files to Existing Files | 8-4 |
| 3. Determine File Systems and Mount Point Directories | 8-5 |

Chapter 8: VHE Configuration and Maintenance (Continued...)

| | |
|--|------|
| 4. Create /etc/vhe_list | 8-6 |
| Example: Simple Configuration with Single File System per Node | 8-7 |
| Example: Node with Multiple File Systems | 8-7 |
| 5. Update /etc/passwd | 8-8 |
| Example | 8-9 |
| Example: Nodes with Multiple File Systems | 8-9 |
| 6. Update /etc/exports | 8-10 |
| 7. Distribute /etc/vhe_list and /etc/passwd | 8-11 |
| 8. Execute /usr/etc/vhe/vhe_mounter | 8-11 |
| Example | 8-12 |
| 9. Verify that VHE is Correctly Configured | 8-13 |
| Configuration Refinements | 8-14 |
| NFS mounts in the Background | 8-14 |
| VHE Maintenance | 8-15 |
| Unmounting file systems | 8-15 |
| Adding or Deleting VHE Nodes | 8-16 |
| Advanced Usage | 8-17 |
| Adding altlogin and mounter Logins | 8-17 |
| Mounter Example | 8-18 |
| Altlogin Example | 8-18 |
| \$HOME | 8-19 |
| \$ROOT | 8-19 |
| Alternate Mount Points | 8-19 |
| Using VHE for Mail | 8-20 |

Chapter 9: Troubleshooting

| | |
|-----------------------------------|-----|
| Key Terms | 9-2 |
| Troubleshooting References | 9-6 |
| Power Up and Connectivity Testing | 9-6 |
| Troubleshooting Sections | 9-7 |
| Guidelines | 9-8 |
| Common Network Problems | 9-8 |
| Initial Troubleshooting | 9-9 |
| Configuration | 9-9 |
| Hardware | 9-9 |

Chapter 9: Troubleshooting (Continued...)

| | |
|--|------|
| Network Communication | 9-10 |
| NFS and Yellow Pages | 9-11 |
| Remote Execution (REX) | 9-12 |
| Error Messages | 9-13 |
| Errnos | 9-14 |
| Unsolved Problems | 9-14 |
| Flowchart Format | 9-15 |
| Troubleshoot NFS | 9-16 |
| Initial Steps to Narrowing the Problem (Flowchart 1) | 9-17 |
| Mount Fails (Flowchart 2) | 9-21 |
| Server Not Responding (Flowchart 3.1) | 9-25 |
| Server Not Responding (Flowchart 3.2) | 9-29 |
| Troubleshoot NFS | 9-31 |
| Restricted Access (Flowchart 4) | 9-33 |
| Programs Hang (Flowchart 5) | 9-37 |
| Performance Problems (Flowchart 6) | 9-41 |
| Troubleshoot Yellow Pages | 9-44 |
| Initial Steps to Troubleshooting YP (Flowchart 7) | 9-45 |
| Incorrect YP Maps (Flowchart 8) | 9-47 |
| ypserv(1M) Problems (Flowchart 9) | 9-51 |
| ypbind(1M) Problems (Flowchart 10) | 9-53 |
| Multiple YP Client Problems (Flowchart 11) | 9-55 |
| Troubleshoot VHE | 9-56 |
| Initial Steps to Troubleshooting VHE (Flowchart 12) | 9-57 |
| Home Node Goes Down After Mount Done (Flowchart 13) | 9-59 |
| Checking /etc/passwd and /etc/vhe_list Files (Flowchart 14) | 9-61 |
| Consistency of /etc/passwd and /etc/vhe_list (Flowchart 15) | 9-63 |
| Execution of vhe_mounter (Flowchart 16) | 9-65 |
| Error Message from vhe_mounter (Flowchart 17) | 9-67 |
| Troubleshoot REX | 9-68 |
| Initial Steps to Troubleshoot REX (Flowchart 18) | 9-69 |
| Initial Steps to Troubleshoot REX (Flowchart 18.1) | 9-71 |
| Unknown Host (Flowchart 19) | 9-73 |

Chapter 9: Troubleshooting (Continued...)

Cannot Connect to REX Server (Flowchart 20) 9-75
User ID Not Valid (Flowchart 21) 9-77
User ID Denied Access (Flowchart 22) 9-79
REX Server Not Running Mount Daemon (Flowchart 23) 9-81
REX Server Denied Access through /etc/exports (Flowchart 24) . . . 9-83
Mount Point Not a Directory (Flowchart 25) 9-87
Command Not Found (Flowchart 26) 9-89
Permission Denied (Flowchart 27) 9-91
Text File Busy (Flowchart 28) 9-93
Device files/named pipes (Flowchart 29) 9-95

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

Appendix B: Migrating from RFA to NFS

Why Migrate to NFS Services? B-1
Similarities B-2
Differences B-2
Changing Scripts from RFA to NFS B-4
Shell Scripts that Accept Different Paths B-4
Shell Scripts with Hard-coded Paths B-5
 Change Pathnames B-5
 Create New Pathnames B-6
RFA through NFS B-7

Appendix C: NFS in an HP-UX Cluster Environment

HP-UX Cluster Terms C-1
NFS Configuration and Maintenance C-2
YP Configuration and Maintenance C-3
Troubleshooting C-3

Appendix D: Password Security

Glossary

Index

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 both the Series 300 and Series 800 HP 9000 computer systems. Any differences in installation, configuration, or operation will be specifically noted.

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, YP (Yellow Pages), VHE (Virtual Home Environment), and REX (Remote Execution Facility) services.
- Routine administration and maintenance of NFS, YP, VHE, or REX.
- Troubleshooting common NFS, YP, VHE, or REX problems

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

Contents

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, YP, 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: YP Configuration and Maintenance

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

The second section explains procedures for administering and maintaining the YP service. It includes topics such as modifying your system to use YP and changing your YP 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, YP, 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: Migrating 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.

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

This manual uses the following format for all entry instructions and examples.

Bold Text emphasizes the word or point.

Computer Text specifies a literal entry. You should enter the text exactly as shown.

Italic Text indicates you should enter information according to your requirements.

Example:

| | | |
|------------------------------|--|--------------------------------------|
| domainname | | <i>domain_name</i> |
| Enter the word domainname | | Enter the name of your YP domain. |

Note Except for the “YP Configuration and Maintenance” chapter, all references to servers and clients apply to NFS servers and clients unless otherwise specified.

Documentation Guide

| For More Information | Read |
|--------------------------------------|--|
| ARPA Services: Daily Use | <i>Using ARPA Services</i> |
| ARPA Services: System Administration | <i>Installing and Administering ARPA Services</i> |
| C Programming Language | <i>C Programming Guide</i> , Jack Purdum, Que Corporation, Indianapolis, Indiana <i>The C Programming Language</i> , Brian W. Kernighan, Dennis M. Ritchie; Prentice-Hall, Inc. |
| Commands and System Calls | <i>ARPA/Berkeley Services Reference Pages</i> <i>HP-UX Reference Manuals</i> <i>NFS Services Reference Pages</i> <i>Network Services Reference Pages</i> |
| HP 92223A Repeater | <i>HP 92223A Repeater Installation Manual</i> |
| HP-UX: Installation | <i>HP-UX Installation Manual</i> |
| HP-UX: Operating System (HP 9000) | <i>HP-UX Concepts and Tutorials</i> <i>HP-UX Installation Manual/HP 9000 Series 300</i> <i>HP-UX Installing and Updating/HP 9000 Series 800</i> <i>HP-UX Reference Manuals</i> <i>HP-UX System Administrator's Manual/HP 9000 Series 800</i> <i>HP-UX System System Administrator's Manual/HP 9000 Series 300</i> <i>Beginner's Guide series for HP-UX</i> <i>Introducing UNIX System V</i> |
| HP-UX: System Administration | <i>HP-UX System Administrator's Manual/HP 9000 Series 800</i> <i>HP-UX System Administrator's Manual/HP 9000 Series 300</i> |

| For More Information | Read |
|--|---|
| LAN Hardware: Installation | <i>HP 98643A LAN/300 Link LANIC Installation Manual LAN Cable and Accessories Installation Manual</i> |
| Networking: General Information | <i>Networking Overview</i> |
| NFS Services: Common Commands | <i>Using NFS Services</i> |
| NFS Services: Programming and Protocols | <i>Programming and Protocols for NFS Services</i> |
| NFS Services: System Administration <ul style="list-style-type: none"> ■ Configuration ■ Installation ■ Maintenance ■ Migrating from RFA to NFS ■ NFS in an HP-UX Cluster Environment ■ NFS Services vs. Local HP-UX ■ Network Lock Manager ■ Remote Execution Facility (REX) ■ Troubleshooting ■ Virtual Home Environment ■ Yellow Pages | <i>Installing and Administering NFS Services</i> |
| NS: System Administration | <i>Installing and Administering NS Services</i> |
| ARPA: System Administration | <i>Installing and Administering ARPA Services</i> |

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.

- 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, YP (Yellow Pages) client, or YP server. (A YP server **must** also be configured as a YP 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, YP client, or YP server. (A YP server **must** also be configured as a YP 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). One advantage is that you can reboot servers without adverse consequences to the client.

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
- Yellow Pages (YP)
- Virtual Home Environment (VHE)

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

Note You must recompile programs that access remote directories. Otherwise, these programs will not be able to access remote directories mounted through NFS since the old directory routines use a *read* call instead of a *getdirentries* call to access those directories.

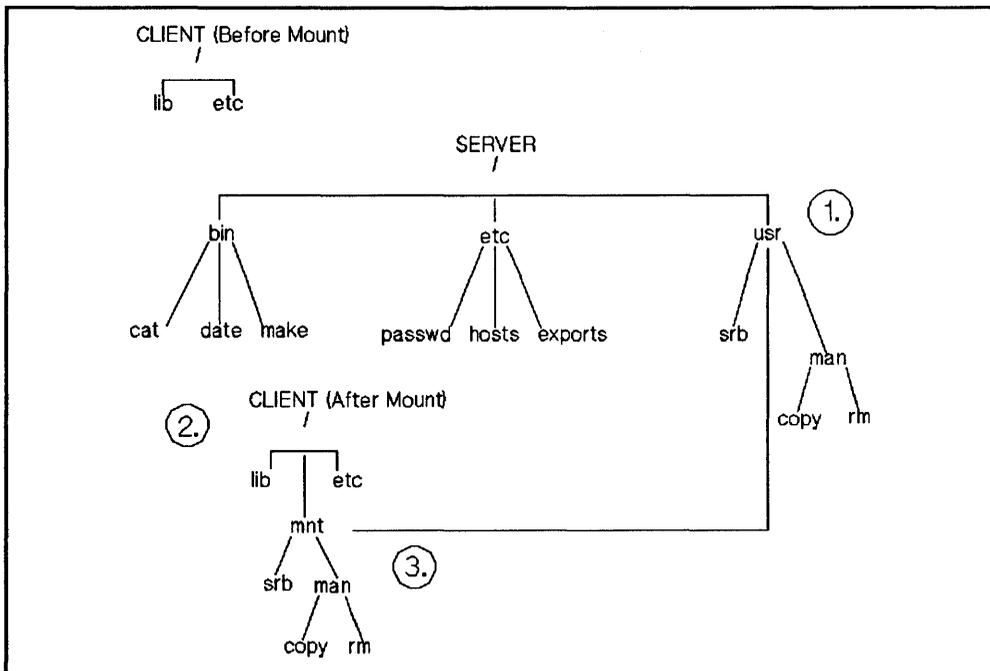
NFS Remote File Access

Before the client can access remote files,

- on the server, the super-user must export the file system (i.e., make it available) to the client and
- on the client, the super-user must mount (import) the file system.

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.

EXAMPLE:



1. The super-user edits the server's */etc/exports* file to make the */usr* file system available to the client.

```
server super-user% cat /etc/exports
/usr client_name
```

2. On the client, the super-user creates a mount point */mnt* (empty directory) and mounts the file system.

```
client super-user% mkdir /mnt
client super-user% 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(2)* 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 super-user 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**.

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*.

EXAMPLE:

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

System C is an NFS client of the NFS server System S, and has mounted file system / from host S on /mnt (super-user 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 with Device Files

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(1m)` command.

EXAMPLE:

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

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

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(1)*) 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 hosts, 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**

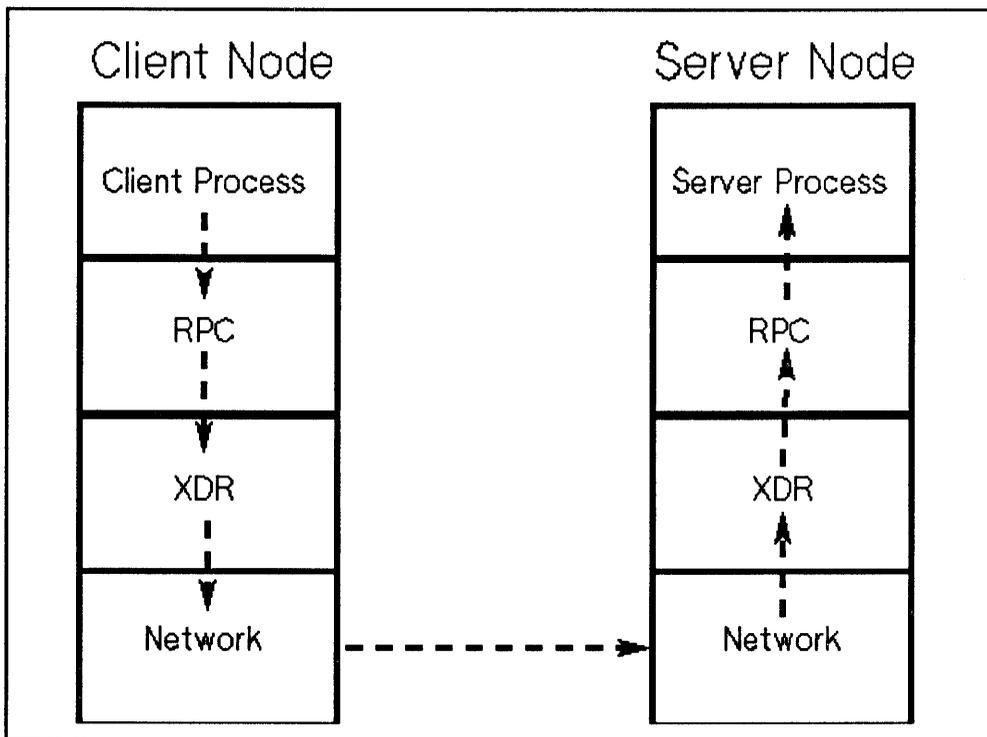
RPCGEN accepts remote program interface definitions (the protocol description file) written in RPC and produces the following C output files, which you may use as a starting point, rewriting as necessary:

- **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 hosts using RPC/XDR. Thus, XDR enables heterogeneous nodes and operating systems to communicate with each other over the network.



RPC and XDR Data Transfer

Note: This figure does not correspond to the ISO Model.

Network Lock Manager

NFS Services includes the Network Lock Manager and the Network Status Monitor. 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.

Yellow Pages (YP)

The Yellow Pages (YP) is an **optional** service containing a collection of cooperating YP server processes that provide YP clients access to data. You can administer all the databases from one YP **master server** since it propagates data across the network to other YP servers. YP includes the following features.

- YP manages unlimited databases. Typically these include files in */etc: group, hosts, netgroup, networks, passwd, protocols, rpc, and 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 YP to obtain information from other YP servers.

- Since the YP master server propagates all **maps** (databases) to the **slave servers**, a YP client receives consistent information regardless of which YP server it accesses.
- If a remote node running a YP server process crashes, YP client processes can obtain YP services from another YP server.
- Since the YP interface uses RPC and XDR, the service is available to other vendors.

YP Advantages

YP has several advantages.

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

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

- YP provides the convenience of centrally administering the */etc* files: *password*, *group*, *hosts*, *netgroup*, *networks*, *rpc*, *services*, and *protocols*.

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

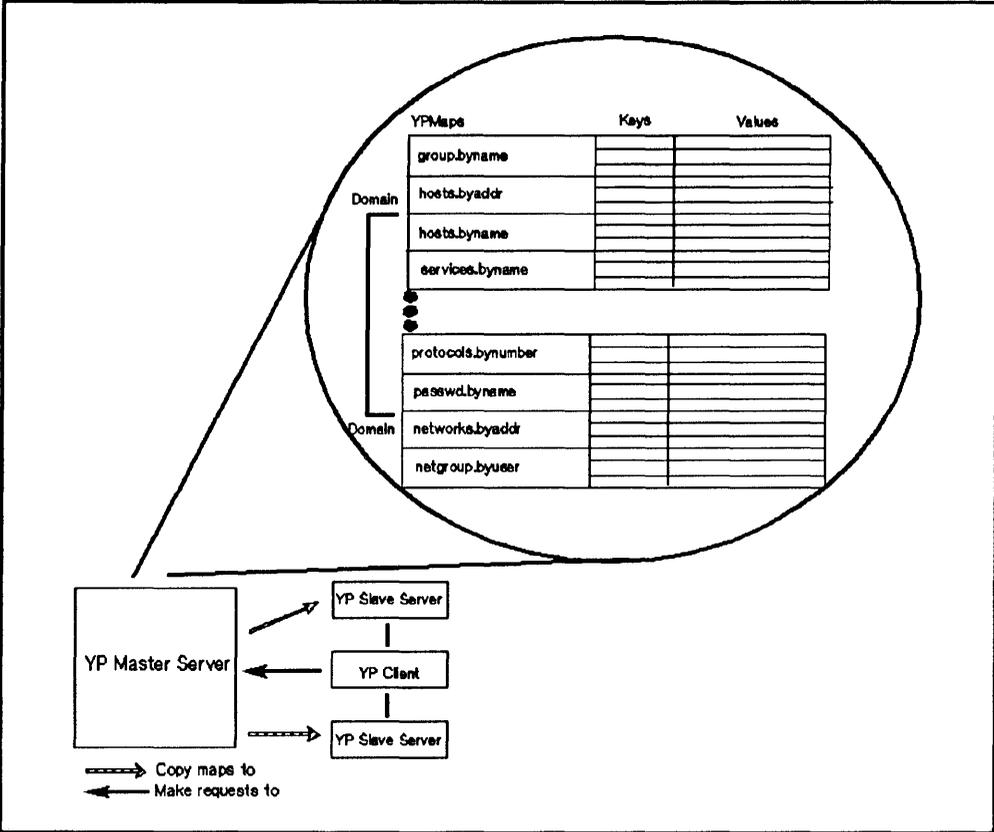
YP Disadvantages

YP has the following disadvantages.

- Since YP provides YP clients access to data via the network, YP clients may observe slower performance than if the data were accessed from local files. For example, with YP, logging in may take more time if the YP server is busy.
- If any of the YP servers are unstable, remote access to files may be slower since the YP client may have to rebind to another YP server. If no other YP server is available, users may not be able to login to their nodes without access to the YP's *passwd* map.
- YP does not make changes visible to all users unless the changes are made on the YP master server.
- The YP slave servers do not immediately see the changes made to the YP master server maps. The updated maps become consistent among all YP servers only after each slave server successfully copies the maps via *ypxfr(1M)*.

YP Concepts

Refer to the following figure and subsections for a summary of how components within Yellow Pages work together: maps, YP domains, YP servers (masters and slaves), and YP clients.



Yellow Pages Structure

YP Maps

The YP system stores information in YP **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.

Note If using YP to provide the information stored in the standard maps' ASCII files, you **must** recompile any applications that read data from those files using standard C library routines.

This recompilation ensures the files can obtain data from the YP maps. If you do not recompile the applications, they will access only the local files. If the local files are not as current as the YP maps, the applications may not work correctly.

YP Servers and YP Clients

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

YP clients are nodes that request access to YP maps from a YP server.

1. A YP client that is not bound sends a broadcast to all YP servers on the network.

2. The YP client **binds** to the first YP server that responds. (Each YP client binds to one YP server per YP domain.)
3. If the request is the YP client's first attempt to access data, the YP client remembers which YP server responded to the request. Subsequent requests by this YP client go directly to this YP server.
4. If the bound YP server is down or unavailable, the YP client automatically rebinds to the first YP server that responds to another broadcast.

Note A YP client can also be configured as any combination of a YP server, NFS client, or NFS server.

A YP server **must** also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.

YP Domains

A **YP domain** is a logical grouping of the set of maps contained on YP servers. You can have different YP domains for multiple sets of nodes on the LAN without worrying about the maps interfering with each other.

- Each one of the nodes within the same YP domain must have the same domain name.
- Maps using the same name in different YP domains can have different contents.

You implement a YP domain as a subdirectory of */usr/etc/yp* on each YP server; the name of this subdirectory is the name of the YP domain. For example, maps in the *research* YP domain would be in */usr/etc/yp/research*. (Note, YP domain names are case sensitive.)

The */etc/netnfsrc* file usually contains the default YP domain name. You can change the default by executing the *domainname(1)* command or by editing */etc/netnfsrc* and then rebooting the system.

YP Masters and YP Slaves

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

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

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

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

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

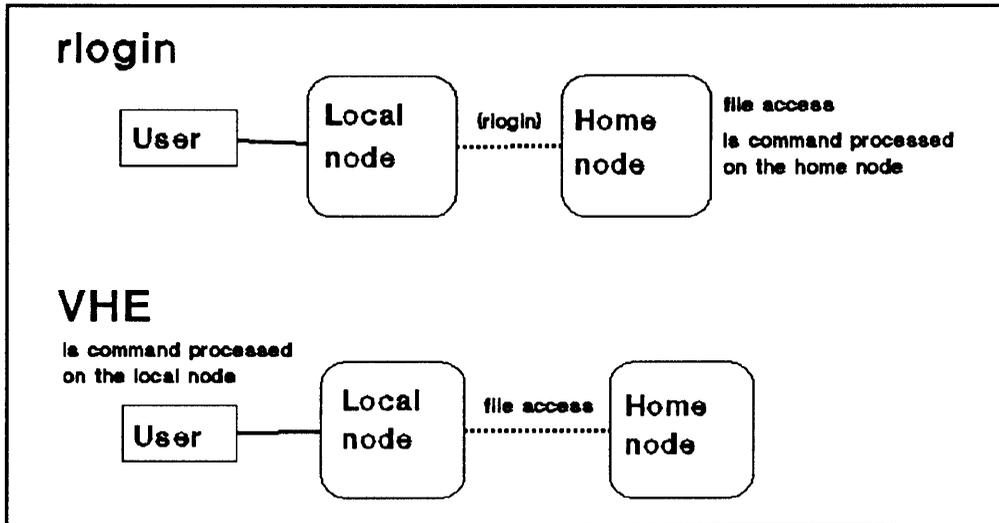
VHE's major advantage is that you can sit down at any remote node (assuming you have login permission), login, and enter into the work environment that is associated with the login on your home node (your home directory as specified in */etc/passwd*). This includes:

- 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.



Comparison: VHE and rlogin Performing an ls command

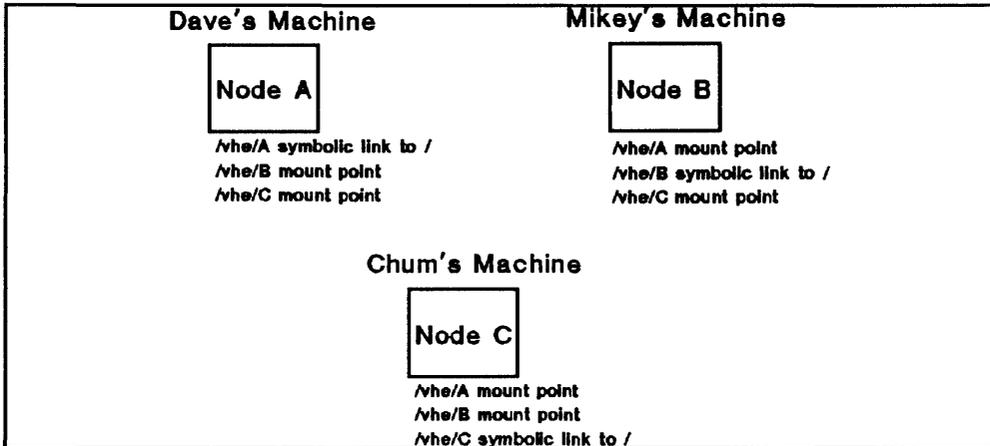
VHE Disadvantages

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 machine. If you try to execute an object code file on the HP 9000 Series 300 machine from the Series 800 machine it will not succeed. However, you can execute a script from the Series 800 machine.
- If you specify pathnames or hardware attributes in your host'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.

How VHE Works

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



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).

Using VHE gives each node access to file systems located on the remote nodes. To maintain consistency when an individual is logged in to his or her 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 users 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 above.

```
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 does 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 and Series 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:
 - Inspect hardware
 - 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 *HP-UX System Administrator's Manual/HP 9000 Series 300* or *HP-UX System Administrator's Manual/HP 9000 Series 800* 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 *HP-UX System Administrator's Manual* for detailed *update* information.
 - 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.
 - 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
5. Relink any of your programs that access remote directories. If you will be using the YP (Yellow Pages) Service, relink programs that call C library routines which access YP files.

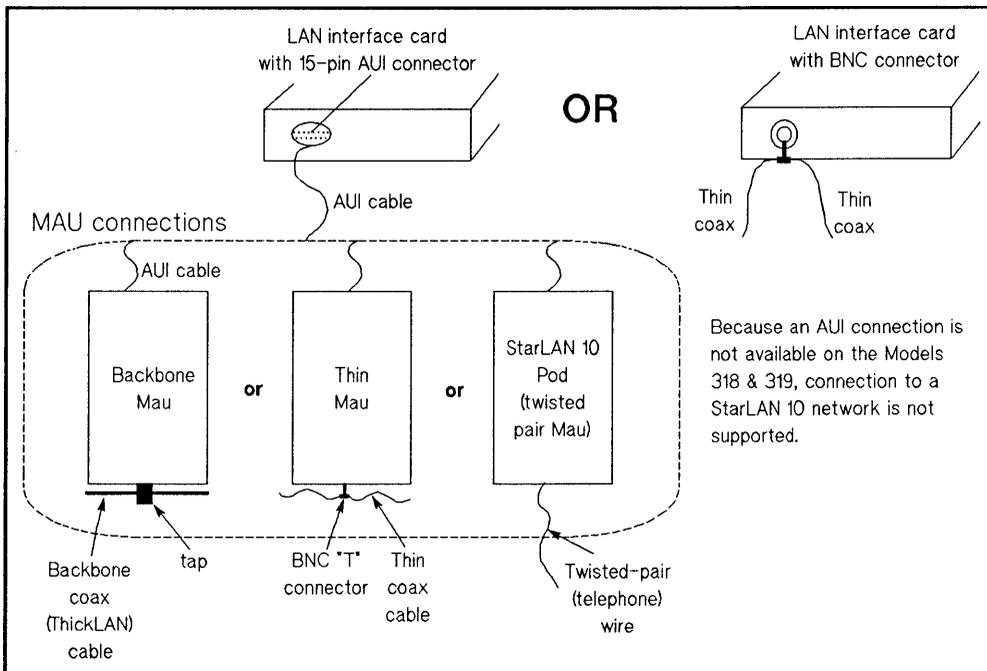
Key Terms

| Term | Definition |
|------------------------------|--|
| CDF | Context Dependent File. A hidden directory that contains all the versions of a file needed by the different cnodes. |
| Cluster | One or more workstations linked together with a local area network (LAN), and sharing a global file system attached to the root server. |
| Cnode | Any node operating in an HP-UX cluster environment, including diskless nodes and the root server. |
| Diskless Cnode | A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a file system directly attached. |
| Heterogeneous Cluster | A diskless cluster with more than one type of computer architecture (e.g., Series 300 and 800) |
| Homogeneous Cluster | A diskless cluster composed of nodes of only one computer architecture (e.g., Series 300 only). |
| 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. |
| NFS | Network File System. |
| Node | A computer system that is attached to or is part of a computer network. |

| Term | Definition |
|--------------------------|--|
| Root Server | The only node in an HP-UX cluster that has file systems directly attached to it. |
| Update | The HP-UX command that installs or updates software onto the system. |
| Yellow Pages (YP) | An optional network service composed of databases (maps) and processes that provide YP clients access to the maps. The YP service enables you to administer these databases from one node. |

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.



Series 300 LAN Interfaces and Cable Connections

HP 9000 Series 300 Models 330, 340, 350, 360 and 370 support any of the LAN hardware connections shown on both sides of the illustration above. These computers are shipped with an installed LAN interface card and your choice of the accompanying hardware for connection to the LAN cable/wire.

HP 9000 Series 300 Models 318 and 319 support the LAN hardware connection shown on the right side of the illustration above. These computers are shipped with an installed interface card that is functionally equivalent to an HP 98643A LAN/300 Link Interface card and an HP 28641A Thin MAU. There is a connector on the backplane of these computers for attaching a T-connector. No MAU or AUI cable is required for these Series 300 models.

HP 9000 Series 800 computers support any of the LAN hardware connections shown on the left side of the illustration above. LAN interface cards for Series 800 computers require installation.

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

- *HP 98643A LAN/300 Link LANIC Installation Manual (Series 300 only)*
- *CIO LAN Interface Controller (LANIC) Installation and Reference Manual or HP Precision Bus Local Area Network Interface Controller (LANIC) Installation and Configuration Guide (Series 800 only)*
- *Twisted-Pair MAU Installation Guide*
- *LAN Cable and Accessories Installation Manual*

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* 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 Series 800 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

Before installing NFS Services, refer to the *HP-UX System Administrator's Manual/HP 9000 Series 300* or *HP-UX System Administrator's Manual/HP 9000 Series 800* 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 super-user 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 NFS Services, Yellow Pages and Virtual Home Environment. These configuration procedures are described later in this manual.

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

For descriptions of the files listed above 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 system, then you have completed installing the NFS Services product. If you have just updated a Series 300 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 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 on a Series 300 (your kernel has not already been configured to include NFS) or
- the *update* program failed to generate a new kernel on a Series 800.

If your kernel is based upon the standard kernel file (*/etc/conf/dfile* on Series 300s or */etc/conf/gen/S800* on Series 800s), 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, your kernel file must contain the uncommented entry

```
nfs
```

On a Series 800, your kernel file must contain the uncommented entry

```
include nfs;
```

Refer to the *HP-UX System Administrator's 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 super-user and refer to the “NFS Configuration and Maintenance,” “YP Configuration and Maintenance,” and “VHE Configuration and Maintenance” chapters to configure your system with NFS, YP (if applicable), and VHE. Refer to the “Configuring and Maintaining the BIND Name Server” chapter in the *Installing and Administering ARPA* manual to configure the BIND Name Server if applicable.

Relink Programs

You **must** relink programs that access remote directories. Otherwise, these programs will not be able to access remote directories mounted through NFS since the old directory routines use a *read* call instead of a *getdirentries* call to access those directories.

If using YP to provide the information stored in the standard maps' ASCII files, you **must** relink any programs that read data from those files using standard C library routines. If you do not relink the programs, they will access only the local files. If the local files are not as current as the YP maps, the programs may not work correctly.

- Programs compiled with previous ARPA/Berkeley libraries *get*ent* (*gethostent*, *getnetent*, *getservent*, *getprotoent* and related routines) will not have access to the YP databases where information in files like */etc/hosts* can be stored.
- Programs compiled with the previous *getpwent* and *getgrent* routines will not have access to information (passwords and groups) stored in the global YP password and group databases.

NFS Configuration and Maintenance

This chapter describes a basic NFS configuration without Yellow Pages. 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

Refer to the “NFS Services Reference Pages” in the *Networking Reference* manual for detailed NFS information.

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 | <ul style="list-style-type: none">■ A node that requests data or services from other nodes (servers).■ A process that requests other processes to perform operations. <p>Note: An NFS client can also be configured as any combination of an NFS server, YP client, or YP server. (A YP server must also be configured as a YP client.)</p> |
| 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. |
| Cnode | Any node operating in an HP-UX cluster environment, including diskless nodes and the root server. |
| Daemon | Background programs that are always running, waiting for a request to perform a task. |
| Diskless Cnode | 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. |
| 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. |
| 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. |

| Term | Definition |
|----------------------------|---|
| 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. |
| 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.) |
| 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 <i>/etc/netgroup</i> . |
| NFS | Network File System. |
| Node | A computer system that is attached to or is part of a computer network. |
| Root Server | The only node in an HP-UX cluster that has file systems directly attached to it. |
| Server | <ul style="list-style-type: none"> ■ A node that provides data or services to other nodes (clients) on the network. ■ A process that performs operations as requested by other processes. <p>Note: An NFS server can also be configured as any combination of an NFS client, YP client, or YP server. (A YP server must also be configured as a YP client.)</p> |
| Soft Mount | An optional mount that causes access to remote file systems to abort requests after one NFS attempt. |

| Term | Definition |
|--------------------------|---|
| 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 <i>mount(1M)</i> command. |
| Update | The HP-UX command that installs software onto the system. |
| Yellow Pages (YP) | <p>An optional network service composed of databases (maps) and processes that provide YP clients access to the maps. The YP service enables you to administer these databases from one node.</p> <p>YP may or may not be active; check with your system administrator.</p> |
| YP Domain | A logical grouping of YP maps (databases) stored in one location. YP domains are specific to the YP network service and are not associated with other network domains. |

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*.

For more information, refer to the *HP-UX System Administrator's Manual/HP 9000 Series 800* or *HP-UX System Administrator's Manual/HP 9000 Series 300* manual.

Configuration Files

The following table lists the files that must be configured for your system to operate correctly. (Refer to the *NFS Services Reference Pages* for detailed information.)

| Configuration File | Description |
|------------------------|--|
| <i>/etc/checklist</i> | Contains a list of file systems that are automatically mounted at boot time. |
| <i>/etc/exports</i> | Contains a list of file systems that clients may import. Note, create this file only on servers. |
| <i>/etc/inetd.conf</i> | Contains information about servers started by inetd(1M), including RPC services. |
| <i>/etc/netgroup</i> | <p>Contains a mapping of network group names (netgroups) to a set of node, user, and YP domain names; both <i>/etc/exports</i> and <i>/etc/passwd</i> can use the netgroups defined in <i>/etc/netgroup</i>.</p> <p>Classifies the nodes for remote mounts.</p> <p>For ARPA Services, classifies the users for remote logins and remote shells. You can specify netgroups in <i>/etc/hosts.equiv</i> and <i>\$HOME/.rhosts</i>.</p> |
| <i>/etc/netnfsrc</i> | Automatically executed at boot time to start the NFS networking (e.g., starts daemons and servers, defines servers and clients). |
| <i>/etc/netnfsrc2</i> | All mounts found in <i>/etc/checklists</i> . |
| <i>/etc/rpc</i> | <p>Maps the RPC program names to the RPC program numbers and vice versa.</p> <p>This file is static; it is already correctly configured.</p> |

| Configuration File | Description |
|---------------------------|---|
| <i>/usr/adm/inetd.sec</i> | 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.

| Daemon | Description |
|-------------------|--|
| <i>biod(1M)</i> | <p>Asynchronous block I/O daemons for NFS clients.</p> <p>If operating in an HP-UX cluster environment, <i>biod(1M)</i> must be running on all cnodes in the cluster.</p> |
| <i>inetd(1M)</i> | <p>Internet daemon that listens on service ports.</p> <ul style="list-style-type: none"> ■ Reads <i>/etc/inetd.conf</i> to determine the appropriate server for handling the incoming request ■ Listens for and accepts network requests ■ Invokes the appropriate server <p>Note: Since <i>inetd(1M)</i> contacts <i>portmap(1M)</i> on behalf of the servers it starts, you must start <i>portmap(1M)</i> before starting <i>inetd(1M)</i>.</p> |
| <i>nfstd(1M)</i> | <p>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 <i>nfstd(1M)</i> process.</p> <p>If operating in an HP-UX cluster environment, <i>nfstd(1M)</i> should be running on the root server if it is servicing NFS requests. Any <i>nfstd(1M)</i> daemons running on client cnodes are ignored.</p> |
| <i>pcnfsd(1M)</i> | <p>Daemon that authenticates a PC user's access to files. It takes the user name and password, and then either</p> <ul style="list-style-type: none"> ■ succeeds (returns a valid UID and GID), or ■ fails (indicates the name and password are unacceptable). <p>Note: Though <i>pcnfsd</i> 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.</p> |

| Daemon | Description |
|--------------------|--|
| <i>portmap(1M)</i> | <p>Daemon that converts RPC program numbers into port numbers. When <i>inetd(1M)</i> starts, it tells <i>portmap(1M)</i>:</p> <ul style="list-style-type: none"> ■ which RPC servers it is listening for, ■ on which ports it is listening, and ■ the RPC program numbers and versions it serves. <p>When a client makes an RPC call to a given program number, it first contacts <i>portmap(1M)</i> on the server node to determine the port number where RPC requests should be sent.</p> <p>Note: Since <i>inetd(1M)</i> contacts <i>portmap(1M)</i> on behalf of the servers it starts, you must start <i>portmap(1M)</i> before starting <i>inetd(1M)</i>.</p> |

Servers

The following table lists the networking servers (processes that perform operations as requested by other processes).

| Server | Description |
|--------------------|---|
| <i>mountd(1M)</i> | <p>Answers file system mount requests by reading <i>/etc/exports</i> to determine which file systems are available to nodes and users; invoked by <i>inetd(1M)</i>.</p> <p>The <i>showmount(1M)</i> command calls <i>rpc.mountd</i> to list the clients with local file systems mounted.</p> <p>If operating in an HP-UX cluster environment, <i>mountd(1M)</i> should be running on the root server if servicing NFS requests. Any <i>mountd(1M)</i> server running on a client cnode is ignored.</p> |
| <i>rstatd(1M)</i> | <p>Returns statistics obtained from the kernel; invoked by <i>inetd(1M)</i>.</p> <p>The <i>rup(1)</i> program uses <i>rpc.rstatd</i>.</p> |
| <i>rusersd(1M)</i> | <p>Lists the users on the local host; invoked by <i>inetd(1M)</i>.</p> <p>The <i>rpc.rusersd</i> server provides the <i>rusers(1)</i> program information about the local users. The <i>rusers(1)</i> program then sums and displays the information.</p> |
| <i>rwalld(1M)</i> | <p>Handles all <i>rwall(1M)</i> requests; invoked by <i>inetd(1M)</i>.</p> <p>The RPC program <i>rwall(1M)</i> sends a message to <i>rpc.rwalld</i> on a given host. Each <i>rpc.rwalld</i> accepts this message and writes it to all users on the host it is serving using <i>wall(1M)</i>.</p> |
| <i>sprayd(1M)</i> | <p>Records the packets sent by <i>spray(1M)</i>; invoked by <i>inetd(1M)</i>.</p> |

NFS Configuration

Configuring your system is the process of setting up your software so it operates correctly and according to your specifications. The following list is an overview of the steps you must complete to configure the nodes on your network with NFS Services. The steps are described in more detail after the overview list.

1. Compare the files in the */etc/newconfig* directory to their corresponding existing files.
2. Set UIDs and GIDs
3. Create an NFS server
 - a. Edit */etc/netnfsrc*
 - b. Edit */etc/inetd.conf*
 - c. Edit */usr/adm/inetd.sec* (if necessary)
 - d. Edit */etc/hosts*
 - e. Edit */etc/netgroup*
 - f. Create and Edit */etc/exports*
 - g. Reboot the system (if necessary)
4. Create an NFS client
 - a. Edit */etc/netnfsrc*
 - b. Mount file systems
 - c. Reboot the system (if necessary)
5. If applicable, configure Yellow Pages (YP).
(Refer to the “YP Configuration and Maintenance” chapter.)
6. If applicable, configure the Virtual Home Environment (VHE) service.
(Refer to “VHE Configuration and Maintenance” chapter.)

7. Execute `/etc/netnfsrc` (or reboot) when you are finished with all of the configuration, including setting up YP and VHE.

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 the NFS service.

- a. Compare each */etc/newconfig* file listed below with its counterpart shown in the following list.

| File in <i>/etc/newconfig</i> directory | Counterpart in <i>/etc</i> directory |
|---|--------------------------------------|
| <i>netgroup</i> | <i>netgroup</i> |
| <i>netnfsrc</i> | <i>netnfsrc</i> |
| <i>rpc</i> | <i>rpc</i> |
| <i>netnfsrc2</i> | <i>netnfsrc2</i> |

- b. If the files are the same, then skip to the next section, “2. Set UIDs and GIDs.”
- c. 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*.

OR

- copy the files in */etc/newconfig* to */etc*. Then customize the files in */etc* if necessary.

2. 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 using Yellow Pages (YP) Service, you can configure YP so you can centrally administer */etc/passwd* and */etc/group*. Note, local UIDs and GIDs are not required if using YP.

If not using YP, 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/Berkeley 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 super-user account would have super-user 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.

3. Create an NFS Server

You must be super-user to create an NFS server.

To create an NFS server, complete the following steps.

- a. Edit `/etc/netnfsrc`
- b. Edit `/etc/inetd.conf`
- c. Edit `/usr/adm/inetd.sec` (if necessary)
- d. Edit `/etc/hosts`
- e. Edit `/etc/netgroup`
- f. Create and Edit `/etc/exports`
- g. Reboot the system (if necessary)

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

Note If operating in an HP-UX cluster environment and configuring NFS on the root server, you must also configure NFS on all clients in the cluster. If the root server does not have NFS configured, then none of the clients can have NFS configured.

A. 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 “4. Create an NFS Client” 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 | <code>NFS_CLIENT=1</code> <code>NFS_SERVER=0</code> |
| Server Only | <code>NFS_CLIENT=0</code> <code>NFS_SERVER=1</code> |
| Both Client and Server | <code>NFS_CLIENT=1</code> <code>NFS_SERVER=1</code> |
| Neither Client nor Server | <code>NFS_CLIENT=0</code> <code>NFS_SERVER=0</code> |
| PC-NFS Server | <code>PCNFS_SERVER=1</code> |

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:
#*     YP specific:
#           domainname      the YP 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
#*     YP specific:
#           ypbind          YP client process (all YP nodes)
#           ypserv          YP server process (YP server only)
#           yppasswdd       YP password daemon (YP 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 (to set the variable LDISK).
##
NFS_CLIENT=0
NFS_SERVER=0
.
.
.
PCNFS_SERVER=0

```

B. 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(1M)` as follows.

```
inetd -c
```

RPC Services Security

The `inetd(1M)` security facility works only when the `inetd(1M)` executes a server. For the RPC services that do not exit after each service request, `inetd(1M)` provides a security check only for the first request. Successive requests bypass the `inetd(1M)` and are subject only to the security checking performed by the individual RPC services. However, you can make the `inetd(1M)` perform a security check for every RPC request by

- adding the `-e` option to the `/etc/inetd.conf` entry for the RPC service and
- then specifying the RPC service in the first field of `/usr/adm/inetd.sec`. (Refer to the next section, “3. 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(1M)* 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(1M)* 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(1)* 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(1)* 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(1)* 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(1M)* 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(1M)* 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.

C. 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 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 and
- which hosts are allowed to remotely use the local host.

Note If `inetd(1M)` 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(1M)`.

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.

service_name allow/deny *host_specifier(s)*

(Note: Enter only *allow* or *deny*)

| /usr/adm/inetd.sec Entry Fields | Description |
|--|---|
| <i>service_name</i> | <p>Name of a valid service (including RPC services) with an entry in <i>/etc/inetd.conf</i>.</p> <ul style="list-style-type: none"> ■ For RPC services, <i>service_name</i> is the name of the service that matches its program number in <i>/etc/rpc</i>. <p>This entry must have a corresponding entry in <i>/etc/inetd.conf</i> which contains the <i>-e</i> option.</p> <ul style="list-style-type: none"> ■ Specify only one service per entry. ■ If an entry in <i>/usr/adm/inetd.sec</i> specifies the service name and nothing else, <i>inetd(1M)</i> allows all hosts to attempt access. |
| <i>allow/deny</i> | <p>The <i>allow</i> entry instructs <i>inetd(1M)</i> to approve the host or network for access to the specified service.</p> <p>The <i>deny</i> entry instructs <i>inetd(1M)</i> to disapprove the host or network for access to the specified service.</p> |
| <i>host_specifier(s)</i> | <p>Name of a host or a network listed in <i>/etc/hosts</i> or <i>/etc/networks</i>, or an internet address in the standard internet notation.</p> <ul style="list-style-type: none"> ■ You can specify more than one host or network by separating each <i>host_specifier</i> 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(1M)* perform its *inetd.sec(4)* 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(4)`.)

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`.

| <code>/usr/adm/inetd.sec</code> Example RPC Entry | Effect on System Security |
|--|---|
| <code>mountd allow hostA</code> | Allows only <i>hostA</i> to access <i>rpc.mountd</i> |
| <code>walld deny 111.56.78.9 10.*</code> | Denies access to <i>rpc.rwalld</i> from the following hosts <ul style="list-style-type: none"> ■ 111.56.78.9 (internet address) ■ all hosts that are part of network 10.* |

D. Edit `/etc/hosts`

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

Each file system in `/etc/exports` may be followed by a netgroup name or a host name. If you want to export local file systems to a specific host, the host name **must** have an entry in `/etc/hosts` regardless of whether it is specified as part of a netgroup in `/etc/netgroup`.

Ensure that a line with the following format exists in `/etc/hosts` for each host to which you will be exporting file systems.

internet_address host_name

If you have configured the BIND name server, see the “Maintaining Network and Domain Data Files” section in the “Configuring and Maintaining The BIND Name Server” chapter in the *Installing and Administering ARPA Services* manual.

- The *internet_address* uniquely identifies the node and must be in decimal dot notation. (Refer to *hosts(4)* for internet address information.)

EXAMPLE: 192.45.36.5 node_7

- The *host_name* is the node name.

The root user should own */etc/hosts* with the permission 0444 (-r—r—r—)

E. Edit */etc/netgroup*

Note If YP is running, do not edit */etc/netgroup* on any node except the YP 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)*.)

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

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

though you must then define these netgroups within `/etc/netgroup`

The `members` is equal to the triple (`host_name`, `user_name`, `yp_domain_name`)

- The entry may contain any number of netgroup names, though you **must** then define these netgroups within `/etc/netgroup`.
- You can assign more than one triple to a netgroup by enclosing each separate set within parentheses (`host_name`, `user_name`, `yp_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` YP 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` YP domain.
- Each `host_name` must have an entry in `/etc/hosts`.
- The `yp_domain_name` is the name of the YP domain to which you currently belong. To determine your current YP domain name, execute the `domainname(1)` command.

The commands using `/etc/netgroup` assume you are not looking for any YP 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 <i>darren</i> on any host in the <i>graphic</i> YP domain |
| netgroup3 (node_7,,graphic) | Any user on the <i>node_7</i> host in the <i>graphic</i> YP domain |
| netgroup4 (node_2,john,) | The user <i>john</i> on the <i>node_2</i> host in any YP domain |
| netgroup5 (,andy,graphic) (node_1,mike,) | The user <i>andy</i> on any host in the <i>graphic</i> YP domain and the user <i>mike</i> on the <i>node_1</i> host in any YP domain |
| netgroup6 (-,annette,graphic) | The user <i>annette</i> in the <i>graphic</i> YP domain, no host included |

F. 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, *mountd(1M)* 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 (see *exports(4)*) for a file system, asynchronous writes on the NFS server occur.

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(2)* or *open(2)* calls),
 - data that cannot be reconstructed (e.g., a file system containing database files),
 - files synchronized with *fsync(2)*, or
 - critical applications requiring absolute data integrity.
 - If you are unsure whether any of the previous conditions apply, do not use the *-async* option.
-

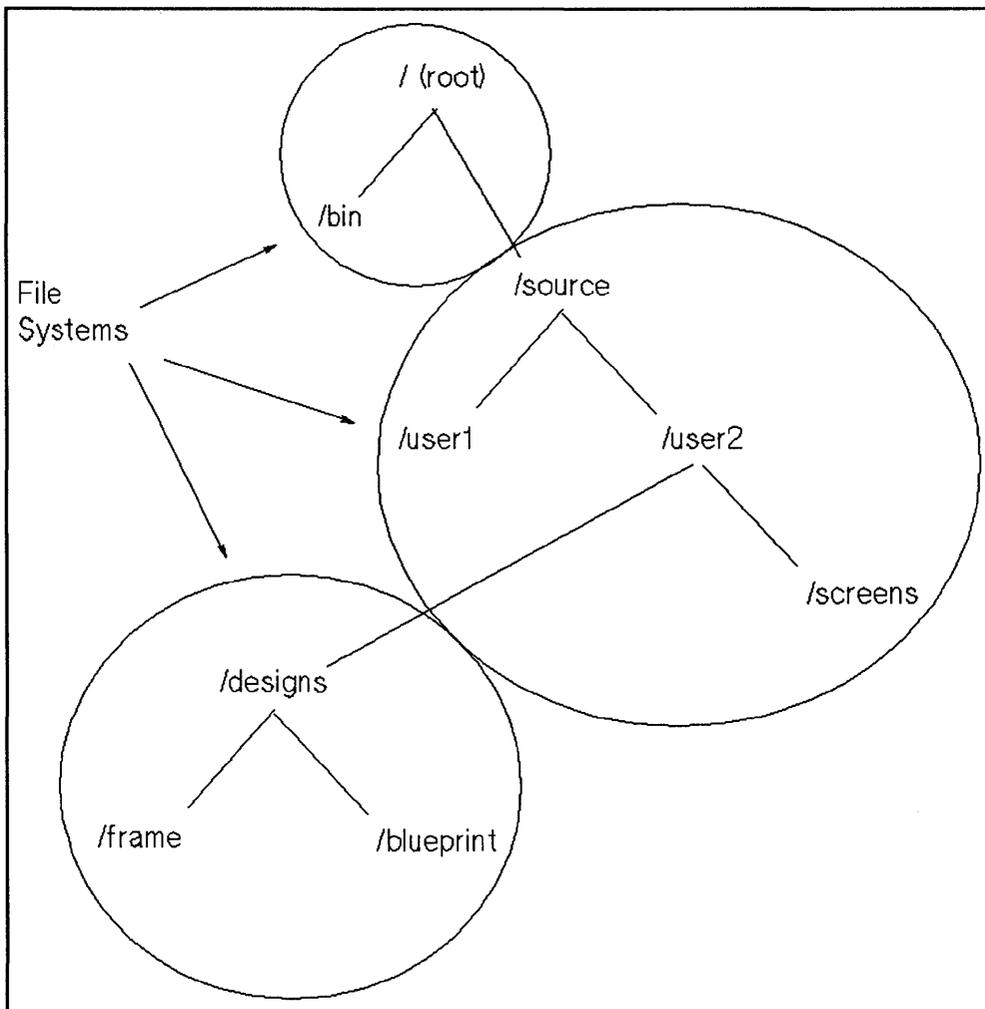
You control the file system's availability by specifying a netgroup or host name; otherwise, the file system becomes available to everyone on the network running NFS. After accessing */etc/exports*, the system checks */etc/netgroup* for the netgroup definition; if it is not present, the system checks */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.

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

EXAMPLE:



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

G. 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.

4. Create an NFS Client

You must be super-user to create an NFS client.

To create an NFS client, complete the following steps.

- a. Edit */etc/netnfsrc*
- b. Mount file systems
- c. Reboot the system (if necessary).

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

a. Edit */etc/netnfsrc*

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

- 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 | <i>NFS_CLIENT</i> =1 <i>NFS_SERVER</i> =0 |
| Server Only | <i>NFS_CLIENT</i> =0 <i>NFS_SERVER</i> =1 |
| Both Client and Server | <i>NFS_CLIENT</i> =1 <i>NFS_SERVER</i> =1 |
| Neither Client nor Server | <i>NFS_CLIENT</i> =0 <i>NFS_SERVER</i> =0 |
| PC-NFS Server | <i>PCNFS_SERVER</i> =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:
#*
#      domainname      the YP 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
#*
#      YP specific:
#      ypbind      YP client process (all YP nodes)
#      ypserv      YP server process (YP server only)
#      yppasswdd      YP password daemon (YP 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 (to set the variable LDISK).
##
NFS_CLIENT=0
NFS_SERVER=0
.
.
.
.
PCNFS_SERVER=0

```

b. 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 if you want it

- mounted automatically at boot time via */etc/checklist* or
- mounted only when manually specified via the *mount(1M)* 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 NS Services* and *Installing and Administering ARPA Services* documentation for more error log information.

Note, if a server goes down from which you previously performed a hard mount, 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(1M)* command. For more specific information, refer to *checklist(4)* and *mount(1M)*.

- 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(1M)* 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(1M)* 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 root server host's name as the node name specified in the *mount(1M)* command.
 - If a cnode mounts a remote file system, any cnode in that cluster can unmount the 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(1M)*.
 - 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 |
|--------------------|---|
| <i>bg</i> | Background: If the first request to a remote node's <i>mountd(1M)</i> fails, the mount process continues retrying the request in the background. |
| <i>defaults</i> | 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 <i>/etc/checklist</i> ; the <i>mount(1M)</i> command automatically provides the defaults. |
| <i>fg*</i> | Foreground: If the first request to a remote node's <i>mountd(1M)</i> fails, the <i>mountd(1M)</i> daemon retries the requests in the foreground. |
| <i>hard*</i> | Hard Mount: NFS retries until the request succeeds or you reboot the system. If using the <i>int</i> default option, you can interrupt the file system request. |
| <i>int*</i> | 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.) |
| <i>noauto</i> | No Automatic Mount: Prevents the file system from being mounted when the <i>mount -a</i> option is executed. You only need to specify <i>noauto</i> when mounting via <i>/etc/checklist</i> . |
| <i>nointr</i> | No Interruptable Mount: You cannot interrupt processes waiting for NFS requests to complete. |
| <i>nosuid</i> | No <i>setuid</i> : You cannot execute files on the remote file system with either the <i>setuid</i> or <i>setgid</i> bits set. |
| <i>port = n</i> | <p>Port = n</p> <p>Default n = 2049 (the NFS server port)</p> <p>Specifies the UDP port at which the NFS server is contacted. You should not have to reset this value.</p> |
| * = Default | |

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

| NFS Mount Options | Description |
|--------------------|--|
| <i>timeo = n</i> | <p>Timeout = n</p> <p>Default n = 7</p> <p>Specifies the initial timeout (in tenths of seconds) for NFS requests.</p> <p>When an NFS request occurs, RPC sends the request, waits 0.7 seconds for a response, and then retries the request.</p> <p>After the initial timeout, the timeout increases by multiples of two each time no response is received. When a specified number of <i>retrans</i> retransmissions have been sent with no reply, soft mounts return an error and hard mounts retry the request.</p> <p>Note: If performing NFS mounts through a gateway and you see several <i>server not responding</i> messages within a few minutes, change the timeout default value (7) to a value of 10 or greater until you stop seeing the message.</p> |
| <i>wsize = n</i> | <p>Write size = n</p> <p>Default n = 8192 bytes (8K)</p> <p>Specifies the maximum write request size used in communicating with the server.</p> |
| * = 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/netnfs2* 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.

NFS Hard Mounts via /etc/checklist

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

OR

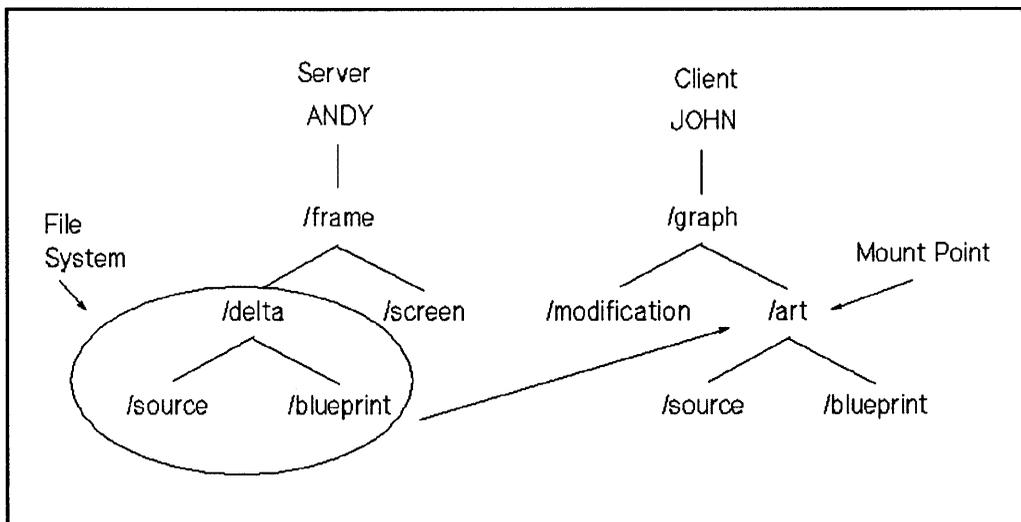
```
server_name:/imported_filesystem /mount_point nfs [options] 0 0
```

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.

EXAMPLES: /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,wsiz=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(1M) for Manual Mounts

Execute *mount(1M)* to mount an NFS file system manually. NFS file systems mounted via *mount(1M)* are only mounted as long as the client is running or until they are unmounted via *umount(1M)*. If the client goes down, you will have to re-mount the file system.

Do not use *mount(1M)* if you listed the file system in */etc/checklist* since it will have already been mounted.

Use the following *mount(1M)* format for NFS mounts. All of the default options are activated unless you specify otherwise.

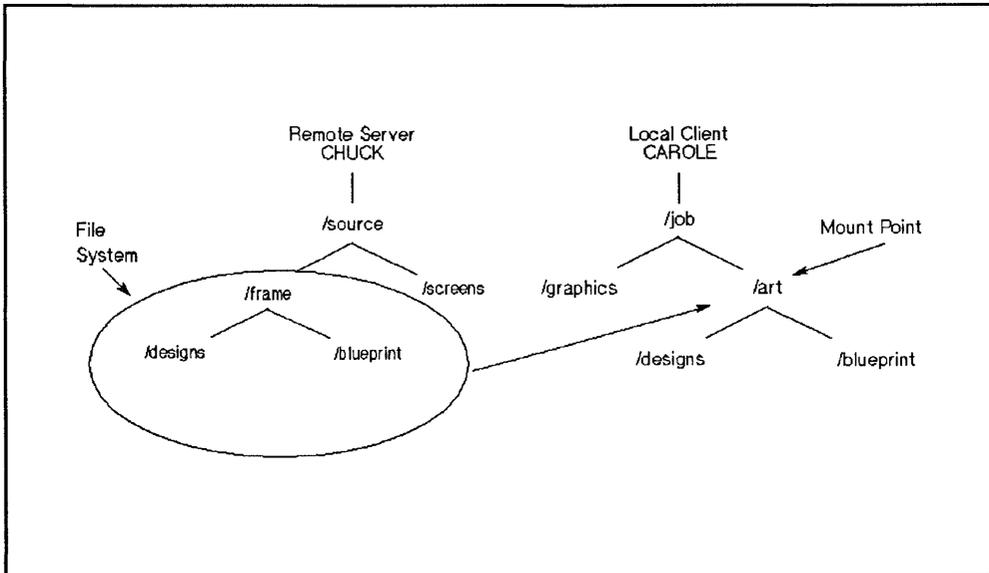
NFS Hard Mounts via mount(1M)

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

NFS Soft Mounts via mount(1M)

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

EXAMPLES: *mount(1M)* Manual NFS Mounts



| mount(1M) 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. | |

c. 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*).

5. Configure YP (optional)

If you plan to use the optional Yellow Pages (YP) service, refer to the “YP Configuration and Maintenance” chapter for detailed configuration procedures.

6. 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.

7. 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.

Maintenance

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

- Remove NFS File Access
- Update Software
- Clock Skew

Remove NFS File Access

You may need to remove file access to NFS either from the:

- client to keep local users from accessing mounted remote file systems, or from the
- server to keep all clients from accessing 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.

- Executing the *umount(1M)* command unmounts file systems mounted either via *mount(1M)* 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 | <p>On the client, execute <i>umount</i>.</p> <p><i>umount mount_point_name</i></p> |
| All file systems on one client | <p>On the client, execute <i>umount -a</i>.</p> <p><i>umount -a</i></p> <p>Note: This command unmounts all file systems, not just NFS file systems.</p> <p>If operating in an HP-UX cluster environment, clients should not execute <i>umount -a</i>.</p> |
| All NFS file systems on all clients | <p>On all clients, execute <i>umount -at</i>.</p> <p><i>umount -at nfs</i></p> |
| All file systems listed in <i>/etc/mnttab</i> that were remotely mounted from a specified server | <p>On all clients, execute <i>umount -h</i>.</p> <p><i>umount -h server_name</i></p> |

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 | <ol style="list-style-type: none"><li data-bbox="676 422 1149 801">1. You have two options for step #1.<ul style="list-style-type: none"><li data-bbox="676 487 1149 619">■ If a netgroup is specified for that file system in <i>/etc/exports</i>, remove the host name from the netgroup entry in the server's <i>/etc/netgroup</i> file.<li data-bbox="676 644 1149 768">■ If a host name is specified for that file system in <i>/etc/exports</i>, remove the host name from the server's <i>/etc/exports</i> file.<li data-bbox="676 801 1149 897">2. On the client, execute <i>umount</i>. <i>umount mount_point_name</i> |
| One NFS file system from a netgroup | <ol style="list-style-type: none"><li data-bbox="676 913 1149 1082">1. On the server, remove the netgroup name (associated with that file system) from either the <i>/etc/exports</i> file or from <i>/etc/netgroup</i>.<li data-bbox="676 1098 1149 1230">2. On all members in the netgroup, execute <i>umount</i>. <i>umount mount_point_name</i> |

| Prevent Access to Server File Systems | Action |
|---------------------------------------|---|
| All NFS file systems from all clients | <p>1. On all clients, execute <i>umount</i>.</p> <p style="padding-left: 40px;"><i>umount mount_point_name</i></p> <p>2. On the server, you have two options for step #2.</p> <p>Kill the <i>nfsd(1M)</i> daemon or daemons (usually four); the system prohibits NFS accesses only until you restart the <i>nfsd(1M)</i> daemons or you reboot the system.</p> <p>Edit <i>/etc/netnfsrc</i> to change the <i>NFS_SERVER=</i> value to zero, and reboot the system.</p> <p><i>NFS_SERVER=0</i></p> |

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/HP 9000 Series 800* or to the *HP-UX System Administrator's Manual/HP 9000 Series 300* 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/inetd.conf*
- */etc/netgroup*
- */etc/netnfsrc2*
- */etc/netnfsrc*
- */etc/rpc*
- */usr/adm/inetd.sec*

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

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 might 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(2)` 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.)

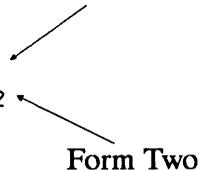
The *ls -l* has two basic forms of output, depending on how old the file is.

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

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

Form One

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 2 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) |
| <hr/> | <hr/> |
| large unsigned number which 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-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 15:30 file3
```


Remote Execution Facility (REX)

Introduction

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

REX consists of:

- The *on(1)* command
- The *rex(1M)* (remote execution daemon)

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

The functionality of REX is similar to that of remote shell (*remsh(1)*) with two important differences:

1. 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.

2. 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 *rex*d 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.

EXAMPLE:

```
on -i -n host
```

is not permitted.

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(1)*, *cs(1)*, and *more(1)*. If this option is specified with a non-interactive command such as *sort(1)*, 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. If your current working directory is

being accessed via RFA, the *on* command will fail with the following error message:

```
on: current directory (<current_dir>) is remote via RFA
RFA directories not supported by on.
```

current_dir is the path name of your current working directory.

- Node B must have *rex*d configured to execute.

Environment Simulation

As mentioned above, 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.
- Finally, the *on* command will fail if your current working directory is being accessed by RFA. This occurs because REX is unable to simulate your environment. In this case, you will receive the following error message:

```
on:      current directory (<current_dir>) is remote via RFA
        RFA directories not supported by on.
```

Configuring *rex*d

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

When *rex*d is configured, it is started by *inetd*(1M) when a request for remote execution is made by a client. *inetd* obtains the information it needs to start *rex*d from the file */etc/inetd.conf*. The following entry must be in the file */etc/inetd.conf* in order for *inetd* to start *rex*d:

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

Where:

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

options are the options that change the behavior of *rex*d. Each of the possible options is described below:

The *-l* option

You can log any errors reported by *rex*d 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, *rex*d appends messages to the file. If *log_file* does not exist, *rex*d 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 *rexid* to use stronger security checking than it uses by default (see “Security Considerations”). When started with the *-r* option, *rexid* 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:

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(1)* 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
```

```
NODE1 mjk
```

Security Considerations

The design and implementation of REX incorporates several security limitations that you should consider before configuring *rex*.

REX restricts access to a system by use of UIDs. That is, the client (*on*) passes the invoking user's UID to the server (*rex*) 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 the *inetd.sec(4)* reference page.

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

Under normal NFS use, only root is allowed to mount remote file systems. However, when *rex* 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 *rex* 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, *rex* will be unable to unmount the file system. Hence, the user has mounted the file system.

Diagnostics

The on Command Error Messages

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 *rex(1M)*.

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: current directory (*<current_dir>*) is remote via RFA
RFA directories not supported by on.

Your current working directory (*<current_dir>*) is being accessed from a remote host via RFA, an HP networking product. Use of *on* from a directory accessed in this manner is not supported.

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 *rex(1M)* found in the *NFS Services Reference Pages*.

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(2)* your current working file system. See *rexd(1M)* in the *NFS Services Reference Pages* 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.

rexid: <host> is not running mountd

The host <host> on which the user's current working directory is located is not running *mountd(1M)*. Therefore, *rexid* is unable to mount the required file system.

rexid: 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, *rexid* is unable to mount the required file system.

The Network Lock Manager

Introduction

This chapter explains file and record locking using the Network Lock Manager and the Network Status Monitor. 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:

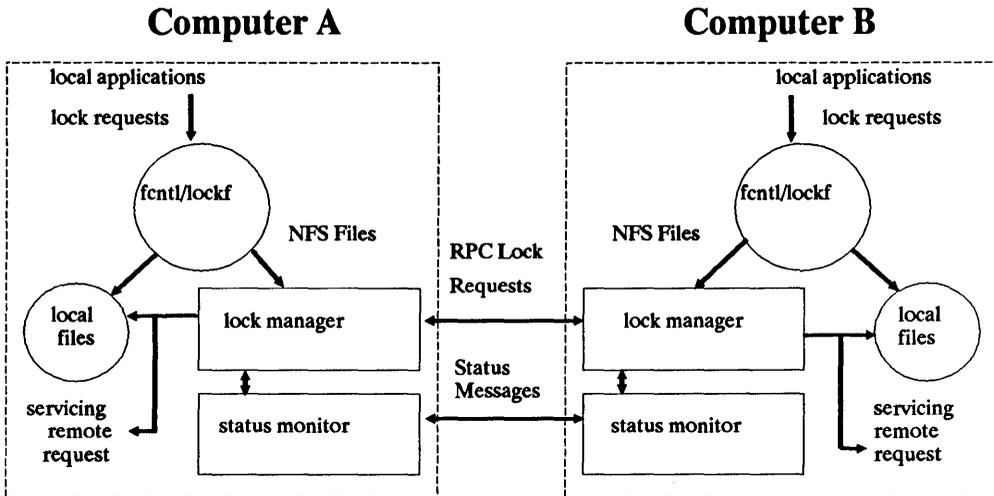
- Basic file access operations, such as read and write, use a stateless protocol (the NFS protocol). All interactions between NFS servers and clients are atomic—the server doesn't remember anything about its clients from one interaction to the next. In the case of a server failure, client applications 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.

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



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, or on behalf of local processes by the kernel. 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 Recovery cannot occur until the remote system is rebooted.

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)* reference pages 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, and the *lockds* release their locks, on the assumption that the lock managers will request locks again when they want 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.

YP Configuration and Maintenance

The Yellow Pages (YP) is an **optional**, distributed network lookup service that allows you to administer databases from one node on the network. With YP you can maintain a single set of user and group IDs for all nodes within a specified set (YP domain). For specific YP information, refer to the following sections:

- Key Terms
- YP Databases
- YP Commands
- YP Configuration
- Verify YP
- Disable YP
- YP Maintenance

Refer to *ypfiles(4)* for a complete explanation of the YP database and directory structure.

If you do not have YP administrative responsibilities, refer to the *Using NFS* manual for general YP usage information.

Note If you configured the BIND name server, it will be used instead of YP for hostname and address resolution. However, YP 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 YP specific.

Key Terms

| Key Terms | Definition |
|-----------------------------|--|
| Bind | <ul style="list-style-type: none"> ■ 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. <p>Acronym for Berkeley Internet Name Domain. The BIND Name Server is a distributed network lookup service.</p> |
| Cluster | One or more workstations linked together with a local area network (LAN), but consisting of only one root file system. |
| Diskless Cnode | A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a file system directly attached. |
| Escape Sequence (YP) | <p>Characters used within files to force inclusion and exclusion of data from YP databases. The escape sequences are as follows.</p> <ul style="list-style-type: none"> ■ + (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. |
| Global (YP) | A means of access in which the system always reads YP 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 Terms | Definition |
|---------------------------|--|
| Key (YP) | 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 <i>passwd.byname</i> map, the users' login names are the keys and the matching lines from <i>/etc/passwd</i> are the values. |
| Local (YP) | A means of access in which the system first reads the local ASCII file. If it encounters an escape sequence, it then accesses the YP databases. |
| Map (YP) | A file consisting of logical records; a search key and related value form each record. YP clients can request the value associated with any key within a map. YP map is synonymous with YP database. |
| Master Server (YP) | The node on which one or more YP maps are constructed from ASCII files. These maps are then copied to the YP slave servers for the YP clients to access. |
| Netgroup | A network-wide group of nodes and users defined in <i>/etc/netgroup</i> . |
| Node | A computer system that is attached to or is part of a computer network. |
| Propagate | To copy maps (data) from one YP server to another. |
| Root Server | The only node in an HP-UX cluster that has file systems directly attached to it. |
| Slave Server (YP) | A node that copies YP maps from the YP master server and then provides YP clients access to these maps. |
| UID | A value that identifies a user in HP-UX. |
| Value (YP) | A unit of information stored in YP maps; each value has a corresponding key (index) so the system can easily retrieve it. For example, in the <i>passwd.byname</i> map, the users' login names are the keys and the matching lines from <i>/etc/passwd</i> are the values. |

| Key Terms | Definition |
|--------------------------|--|
| Yellow Pages (YP) | <p>An optional network service composed of databases (maps) and processes that provide YP clients access to the maps. The YP service enables you to administer these databases from one node.</p> <p>YP may or may not be active; check with your system administrator.</p> |
| YP Client | <ul style="list-style-type: none"> ■ A node that requests data or services from YP servers. ■ A YP process that requests other YP processes to perform operations. <p>Note: A YP client can also be configured as any combination of a YP server, NFS client, or NFS server. (Note, a YP server must also be configured as a YP client.)</p> |
| YP Database | See "Map (YP)." |
| YP Domain | A logical grouping of YP maps (databases) stored in one location. YP domains are specific to the YP network service and are not associated with other network domains. |
| YP Map | See "Map (YP)." |
| YP Password | <p>The password for a user's login ID that exists in the YP <i>passwd</i> map. The YP password is the same one as the user password, but is administered through the YP.</p> <p>You do not have to have a YP password to access the YP databases.</p> |
| YP Server | <ul style="list-style-type: none"> ■ A node that provides data (maps) or services to other nodes (YP clients) on the network using YP. ■ A YP process that performs operations as requested by other YP processes. <p>Note: A YP server must also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.</p> |

YP Databases

The *ypmake(1M)* script creates the standard YP databases from the following ASCII files. You can also create additional YP databases. (Refer to *ypfiles(4)*.)

| | |
|----------------------|-----------------------|
| <i>/etc/group</i> | <i>/etc/passwd</i> |
| <i>/etc/hosts</i> | <i>/etc/protocols</i> |
| <i>/etc/netgroup</i> | <i>/etc/rpc</i> |
| <i>/etc/networks</i> | <i>/etc/services</i> |

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(1M)* command automatically adds the *.dir* and *.pag* suffixes to each map name. Refer to the *HP-UX System Administrator's Manual/HP 9000 Series 800* or *HP-UX System Administration Tasks/HP 9000 Series 300* manual for more information on short file name file systems.

Local and Global Maps

Clients access the above ASCII files and their corresponding YP maps in one of two ways, depending on whether the YP 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 YP database.
- A map is **global** if the system accesses only the YP database (never accesses the local ASCII file).

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

| YP Maps | Type | Access |
|---|--------|---|
| <i>/etc/group</i> <i>/etc/passwd</i> | Local | <p>If a + (plus) entry exists at the beginning of a line, the system retrieves data from the corresponding YP map; otherwise, the YP maps are unused.</p> <p>Occurrences of <i>+</i><i>@netgroup_name</i> and <i>-@netgroup_name</i> at the beginning of a line cause the system to reference YP.</p> <p>(Refer to <i>group(4)</i> and <i>passwd(4)</i> for complete information regarding these escape sequences.)</p> |
| <i>/etc/hosts</i> <i>/etc/netgroup</i> <i>/etc/networks</i> <i>/etc/protocols</i> <i>/etc/rpc</i> <i>/etc/services</i> | Global | <p>The system consults only the YP for data. If YP is not running, it looks at the local ASCII files.</p> <p>However, if the BIND Name Server is configured, the system will use it for host name and address resolution instead of YP.</p> |

Note

If using YP to provide the information stored in the standard maps' ASCII files, you **must** relink any applications that read data from those files.

This relinking ensures the files can obtain data from the YP maps. If you do not relink the applications, they will access only the local files. If the local files are not as current as the YP maps, the applications may not work correctly.

Escape Sequences

Escape sequences are characters used within a file at the beginning of a line to force inclusion and exclusion of data from YP databases. (Refer to the following reference pages: *passwd(4)*, *hosts(4)*, *netgroup(4)*, *host.equiv(4)*, and *group(4)*.) The escape sequences are as follows.

- + (plus)
- - (minus)
- +@netgroup_name
- -@netgroup_name

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

| Escape Sequence | Description |
|------------------------|--|
| <i>-@netgroup_name</i> | <p>Use <i>-@netgroup_name</i> in <i>/etc/passwd</i> to disallow the matching entries from the YP <i>passwd</i> map for all members of a network group.</p> <p>For ARPA Services Use <i>-@netgroup_name</i> in <i>/etc/hosts.equiv</i> and <i>\$HOME/.rhosts</i> to exclude a network group's entries from their lists of allowed users.</p> |

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(1)*) and remote shell execution (*remsh(1)*).

The master server uses */etc/netgroup* to generate three YP 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 programs consulting the YP *netgroup* maps include *login(1)*, *mountd(1M)*, *rlogin(1)*, and *remsh(1)*.

| Program | Description |
|-------------------------------------|--|
| <i>login(1)</i> | Consults the maps to resolve netgroup names in <i>/etc/passwd</i> |
| <i>mountd(1M)</i> | Consults the maps to resolve netgroup names in <i>/etc/exports</i> |
| <i>rlogin(1)</i> <i>remsh(1)</i> | <p>For ARPA Services Consults the <i>netgroup</i> map if netgroup names are in <i>/etc/hosts.equiv</i> or <i>\$HOME/.rhosts</i></p> |

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)*.)

The entry may contain any number of netgroup names,

netgroup_name1 netgroup_name2 netgroup_name3 ...

netgroup_name1 member1 member2 ...

though you must then define these netgroups within

/etc/netgroup

The *members* is equal to the triple (*host_name*, *user_name*, *yp_domain_name*)

- The entry may contain any number of netgroup names, though you **must** then define these netgroups within */etc/netgroup*.
- You can assign more than one triple to a netgroup by enclosing each separate set within parentheses (*host_name*, *user_name*, *yp_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* YP 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* YP domain.
- Each *host_name* must have an entry in */etc/hosts*. (See *hosts(4)*.)

- The *yp_domain_name* is the name of the YP domain to which you currently belong. The commands using */etc/netgroup* assume you are not looking for any YP domain other than the one assigned on your node. (To list your current YP domain name, execute the *domainname(1)* command.)

EXAMPLE: The following example is a sample */etc/netgroup* file. (Refer to *netgroup(4)* 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 YP domain name for all the example netgroups is *mickie*. The users and hosts are classified into netgroups as follows.

| 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 YP <i>passwd</i> map | no hosts |
| allhosts | no users | all hosts in the YP <i>hosts</i> map |

Files Related to YP

For ARPA Services

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

```
+@engineering
```

in */etc/hosts.equiv* will include all members of engineering as defined in the local file */etc/netgroup* or in the YP 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(1)* or *remsh(1)*. (See *hosts.equiv(4)*.)

YP Commands

Refer to the following table for a brief description of all YP commands. Refer to the “Common Commands” chapter in the “Using NFS Services” manual for a more detailed description of the YP commands you might want to use on a daily basis (i.e., those YP commands that do not require super-user access).

| YP Commands | Description |
|----------------------|--|
| <i>domainname(1)</i> | Use <i>domainname(1)</i> to determine or change a YP domain name. |
| <i>makedbm(1M)</i> | <p>Note: Use this version of <i>makedbm(1M)</i> only with YP.</p> <p>A tool for building YP maps.</p> <p>Use <i>makedbm(1M)</i> to build or rebuild databases not built by <i>/usr/etc/yp/ypmake</i>.</p> <p>Use <i>makedbm(1M)</i> to disassemble a map so that you can see the key-value pairs comprising it.</p> |
| <i>ypbind(1M)</i> | <p>Used by each client to determine to which server it should bind.</p> <p>Note: This entry exists in the <i>NFS Services Reference Pages</i> as <i>ypserv(1M)</i>; it exists online as <i>ypbind(1M)</i>.</p> |
| <i>ypcat(1)</i> | Lists the contents of a YP map. |
| <i>ypinit(1M)</i> | <p>On YP master servers, <i>ypinit(1M)</i> constructs maps from <i>/etc</i> files.</p> <p>On YP slave servers, <i>ypinit(1M)</i> copies the initial map versions from the master server.</p> |

| YP Commands | Description |
|----------------------|--|
| <i>yppmake(1M)</i> | A script, initially called by <i>ypinit(1M)</i> , that builds standard YP maps from ASCII files. These files are usually in <i>/etc: passwd, hosts, group, netgroup, networks, protocols, rpc, and services</i> . |
| <i>yppmatch(1)</i> | Prints the value for one or more specified keys in a YP map. |
| <i>yppasswd(1)</i> | Changes the password for your current login ID in the YP <i>passwd</i> map. (You do not have to have a YP password to access the YP databases.) |
| <i>yppasswdd(1M)</i> | A server, running only on the master server, that permits users to change their password in the YP <i>password</i> map. |
| <i>yppoll(1M)</i> | Asks any <i>ypserv(1M)</i> for the information it holds about a single map. |
| <i>yppush(1M)</i> | Used by the master server to administer a running YP system. The <i>yppush(1M)</i> command causes a YP map to be copied (using <i>ypxfr(1M)</i>) from the maps' master server to each slave server in the YP domain. |
| <i>ypserv(1M)</i> | Provides access to data stored in YP maps on servers. If operating in an HP-UX cluster environment, <i>ypserv(1M)</i> should be running on the root server. |
| <i>yppset(1M)</i> | Tells the local <i>ypbind(1M)</i> process to obtain YP services for a YP domain from a specific server. |

| YP Commands | Description |
|-------------------|---|
| <i>ypwhich(1)</i> | Tells you which server a node is currently using or which server is master of a specified map. |
| <i>ypxfr(1M)</i> | Transfers a YP map from one server to another. Run <i>ypxfr(1M)</i> one of three ways: <ul style="list-style-type: none"> ■ <i>yppush(1M)</i> periodically, ■ <i>ypxfr(1M)</i> interactively, or ■ <i>via cron(1M)</i> periodically. |

YP Configuration

Refer to the following checklist for YP configuration steps.

1. Create a YP Master Server
2. Create a YP Client
3. Create a YP Slave Server
4. Propagate the YP Maps

Before Configuring YP

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 YP.

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

| File in /etc/newconfig directory | Counterpart in /usr/etc/yp directory |
|----------------------------------|--------------------------------------|
| <i>ypinit</i> | <i>ypinit</i> |
| <i>yp_Makefile</i> | <i>Makefile</i> |
| <i>ypmake</i> | <i>ypmake</i> |
| <i>ypxfr_1perday</i> | <i>ypxfr_1perday</i> |
| <i>ypxfr_1perhour</i> | <i>ypxfr_1perhour</i> |
| <i>ypxfr_2perday</i> | <i>ypxfr_2perday</i> |

2. If the files are the same, skip to the next section, “1. Create a YP 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*.

OR

- copy the files in */etc/newconfig* to */usr/etc/yp*. Then re-customize the files in */usr/etc/yp* if necessary.

1. Create a YP Master Server

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

A YP server **must** also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.

Before Creating a YP 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*. (Refer to the *NFS Services Reference Pages*.)
2. If you know the correct configuration, create the */etc/netgroup* file. (Refer to *netgroup(4)*.)

Note The YP maps store only the first occurrence if

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

Security

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.yp*).
2. Delete undesired users from the original */etc/passwd* file. To prevent all entries in the YP *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: PWFIL=/etc/passwd
```

```
To: PWFIL=/etc/passwd.yp
```

4. Edit */etc/netnfsrc* as follows.

```
Change: /usr/etc/rpc.yppasswdd /etc/passwd -m passwd PWFIL=/etc/passwd
```

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

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

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

If in the future you need to run *ypmake(1M)* and you have restricted access to the master server as just described, enter the following line.

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

Note For information on C2 Security, refer to the *HP-UX System Security Manual, A Beginner's Guide to Using Shells/HP 9000 Series 300/800 Computers*, and the *HP-UX Beginner's Guide*.

Creating a YP Master Server

Refer to the following steps to create your master server. **If operating in an HP-UX cluster environment**, HP recommends that the root server be the master server.

1. Set the YP domain name using the *domainname(1)* command. This YP domain name must be the same one used for all clients and servers within this YP domain.

```
domainname YP_domain_name
```

2. Execute *ypinit(1M)* with the *-m* parameter.

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

3. The system asks whether you want the procedure to quit at the first non-fatal error. You may want to answer “no” since you can later correct the errors.
 - a. Respond *no* or *n* for *ypinit(1M)* to continue regardless of the errors. After the procedure finishes, correct all errors that occurred.

| Manually Starting YP Master Server | Automatically Starting YP Master Server (at Boot Time) |
|--|--|
| <p>1. If you have not already done so, set the YP domain name using the <i>domainname(1)</i> command. This YP domain name must be the same one used for all clients and servers within this YP domain.</p> <pre>domainname YP_domain_name</pre> <p>2. Execute <i>ypserv(1M)</i>.</p> <pre>/usr/etc/ypserv</pre> <p>Note: If operating in an HP-UX cluster environment, start <i>ypserv(1M)</i> only on the root server, and start <i>ypbind(1M)</i> on every cnode.</p> <p>3. Execute <i>ypbind(1M)</i>.</p> <pre>/etc/ypbind</pre> | <p>Edit <i>/etc/netmfsrc</i> to perform the following actions.</p> <p>Note: A zero in the <i>YP_CLIENT</i>, <i>YP_MASTER_SERVER</i>, or <i>YP_SLAVE_SERVER</i> field disables the node from working as a client, master server, or slave server respectively.</p> <ul style="list-style-type: none"> ■ Set <i>YPDOMAIN</i> to the YP domain name. <pre>YPDOMAIN=YP_domain_name</pre> <p>You will need to use this same YP domain name for all clients and servers within this YP domain.</p> <ul style="list-style-type: none"> ■ Set <i>YP_MASTER_SERVER</i> to a value other than zero. Changing this variable permits users to change their YP password. <pre>YP_MASTER_SERVER=1</pre> <ul style="list-style-type: none"> ■ Set the <i>YP_SLAVE_SERVER</i> to zero to disable the node as a slave server. <pre>YP_SLAVE_SERVER=0</pre> <ul style="list-style-type: none"> ■ Set <i>YP_CLIENT</i> to a value other than zero. <pre>YP_CLIENT= 1</pre> |

2. Create a YP Client

You must be super-user to create a YP client.

A YP client can also be configured as an NFS client, NFS server or both. All YP servers **must** also be configured as YP clients.

Before Creating a YP Client

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

Creating a YP Client

Customize the following files which traditionally store the information. (For suggested modifications, refer to the following “Altering a Client’s Files” section.)

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

| | |
|----------------------|-----------------------|
| <i>/etc/group</i> | <i>/etc/passwd</i> |
| <i>/etc/hosts</i> | <i>/etc/protocols</i> |
| <i>/etc/netgroup</i> | <i>/etc/rpc</i> |
| <i>/etc/networks</i> | <i>/etc/services</i> |

Altering a Client's Files

The following table provides suggestions for altering the client files.

| Client File | Suggested Modification |
|-------------------|--|
| <i>/etc/group</i> | <p>You may want to reduce <i>/etc/group</i> 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 the YP service since group is a local map.</p> <p>+:</p> |
| <i>/etc/hosts</i> | <p>Ensure <i>/etc/hosts</i> contains an entry for the local host name. The system accesses these entries when the YP service is not yet available. After the <i>ypbind(1M)</i> process is running, the system never accesses <i>/etc/hosts</i>.</p> <p>EXAMPLE: Sample YP client's <i>/etc/hosts</i> entry</p> <pre>192.9.1.87 local_host # Byron W. Donne11</pre> |

| Client File | Suggested Modification |
|--|---|
| <p><i>/etc/hosts.equiv</i> For ARPA Services</p> | <p>The system first accesses <i>/etc/hosts.equiv</i> directly. If a <i>+@netgroup</i> or <i>-@netgroup</i> entry exists, the system accesses the YP <i>netgroup</i> map.</p> <p>Note, using <i>netgroup</i> reduces <i>rlogin(1)</i> and <i>remsh(1)</i> problems that occur because different <i>/etc/hosts.equiv</i> files exist on different nodes.</p> <p>For more control over logins, edit <i>/etc/hosts.equiv</i> as follows.</p> <ol style="list-style-type: none"> 1. Enter either a plus (+) or (-) to enable or disable <i>login</i>, respectively. 2. Enter the at (@) character. 3. Enter the <i>netgroup_name</i> as defined in the global <i>netgroup</i> database. <p>EXAMPLE:</p> <pre> +@netgroup1_name (trusted) +@netgroup2_name (trusted) -@netgroup3_name (distrusted) </pre> |
| <p><i>\$HOME/.rhosts</i> For ARPA Services</p> | <p>The system first accesses <i>\$HOME/.rhosts</i> directly. If a <i>+@netgroup</i> or <i>-@netgroup</i> entry exists, the system accesses the YP <i>netgroup</i> map. (Refer to the above <i>/etc/hosts.equiv</i> example.)</p> <p>Since the super-user's <i>\$HOME/.rhosts</i> controls remote super-user access to the local node, HP recommends restricted access. To restrict access, either make the list of trusted hosts explicit or use <i>netgroup</i> names.</p> |

| Client File | Suggested Modification |
|--------------------|---|
| <i>/etc/passwd</i> | <p>Ensure <i>/etc/passwd</i> contains</p> <ul style="list-style-type: none"> ■ entries for the root user, ■ entries for the primary users, and ■ an escape entry to use the YP service. <p>Entries in the local <i>/etc/passwd</i> file mask identical name entries in the YP <i>passwd</i> maps. Delete all other names and enter <code>+::0:0:::</code> as the last line. This line causes library routines looking for a particular entry to search the YP database.</p> <p><code>+::0:0:::</code></p> <p>EXAMPLES: Sample entries in <i>/etc/passwd</i></p> <pre>+ap::::Dave Hamil:/usr2/ap:/bin/csh</pre> <p>The system pulls an entry for <i>ap</i> from the YP <i>passwd</i> map because of the + (plus) escape character.</p> <p>It obtains the <i>UID</i>, <i>GID</i>, and <i>password</i> from the YP and obtains the comment field, home directory, and default shell from the <i>/etc/passwd</i> entry.</p> <p>If no entry for <i>ap</i> exists in the YP, the system reacts as though no entry for <i>ap</i> exists anywhere.</p> <pre>ap::140:100:Mike Donn:/usr2/ap:/bin/csh</pre> <p>Since the plus (+) escape character is not present, the system does not access YP. User <i>ap</i> has no password.</p> |

| Client File | Suggested Modification |
|--------------------|--|
| <i>/etc/passwd</i> | <p data-bbox="427 227 982 257">EXAMPLES: Sample entries in <i>/etc/passwd</i></p> <p data-bbox="427 299 475 320">+ap:</p> <p data-bbox="427 353 1123 417">The system obtains all information from the YP <i>passwd</i> map for user <i>ap</i>.</p> <p data-bbox="427 488 532 510">+::0:0:::</p> <p data-bbox="427 543 1123 607">The system obtains all information from the YP <i>passwd</i> map for all users not already encountered.</p> |

Starting the YP Client

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

| Manually Starting YP Client | Automatically Starting YP Client (at Boot Time) |
|---|---|
| <p>1. If you have not already done so, set the YP domain name using the <i>domainname(1)</i> command. This YP domain name must be the same one used for all clients and servers within this YP domain.</p> <pre>domainname YP_domain_name</pre> <p>2. Execute <i>ybind(1M)</i>.</p> <pre>/etc/ybind</pre> | <p>Edit <i>/etc/netnfsrc</i> to perform the following actions.</p> <ul style="list-style-type: none">■ Set <i>YPDOMAIN</i> to the same YP domain name used on all clients and servers within this YP domain. <pre>YPDOMAIN=YP_domain_name</pre> <ul style="list-style-type: none">■ Set <i>YP_CLIENT</i> to a value other than zero. <pre>YP_CLIENT=1</pre> <p>Note: A zero in the <i>YP_CLIENT</i> field disables the node from working as a YP client.</p> |

Note If you want the node to be a server also, refer to either the “1. Create a YP Master Server” or “3. Create a YP Slave Server” section for complete instructions.

3. Create a YP Slave Server

You must be super-user to create a YP slave server.

You may want to create slave servers to improve the reliability of your system.

A YP server **must** also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.

Before Creating a YP Slave Server

Before creating a slave server, ensure the following:

- master server exists (see “1. Create a YP Master Server” section) and
- *ypserv(IM)* is running on the master server.

Creating a YP Slave Server

Refer to the following steps to create a slave server.

1. Set the YP domain name using the *domainname(1)* command. This YP domain name must be the same one used for all clients and servers within this YP domain.

```
domainname YP_domain_name
```

2. Execute *ypinit(1M)* with the *-s* parameter.

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

3. The system asks whether you want the procedure to quit at the first non-fatal error. You may want to answer “no” since you can later correct the errors.

Note If you want this node to serve a YP domain that is different from the one set by the *domainname(1)* command, use the *DOM=* parameter to *ypinit(1M)*. (For details, refer to *ypinit(1M)*.) If you use the *DOM=* parameter, ensure that the master server serves the domain that you specify.

- a. Respond *no* or *n* for *ypinit(1M)* to continue regardless of the errors. After the procedure finishes, correct all errors that occurred.
 - b. Respond *yes* or *y* for *ypinit(1M)* 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(1M)* 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 a YP Client’s Files” in the “2. Create a YP Client” section.

Starting the YP 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 YP Slave Server | Automatically Starting YP Slave Server (at Boot Time) |
|--|---|
| <p>1. If you have not already done so, set the YP domain name using the <i>domainname(1)</i> command. This YP domain name must be the same one used for all clients and servers within this YP domain.</p> <pre>domainname YP_domain_name</pre> <p>2. Execute <i>ypserv(1M)</i>.</p> <pre>/usr/etc/ypserv</pre> <p>Note: If operating in an HP-UX cluster environment, start <i>ypserv(1M)</i> only on the root server, and start <i>ypbind(1M)</i> on every cnode.</p> <p>3. Execute <i>ypbind(1M)</i>.</p> <pre>/etc/ypbind</pre> | <p>Edit <i>/etc/netnfsrc</i> to perform the following actions.</p> <p>Note: A zero in the <i>YP_CLIENT</i>, <i>YP_MASTER_SERVER</i>, or <i>YP_SLAVE_SERVER</i> field disables the node from working as a client, master server, or slave server respectively.</p> <ul style="list-style-type: none"> ■ Set <i>YPDOMAIN</i> to the same YP domain name used on all clients and servers within this YP domain. <pre>YPDOMAIN=YP_domain_name</pre> <ul style="list-style-type: none"> ■ Set the <i>YP_MASTER_SERVER</i> to zero to disable the node as a master server. <pre>YP_MASTER_SERVER= 0</pre> <ul style="list-style-type: none"> ■ Set <i>YP_SLAVE_SERVER</i> to a value other than zero. <pre>YP_SLAVE_SERVER=1</pre> <p>Set <i>YP_CLIENT</i> to a value other than zero.</p> <pre>YP_CLIENT=1</pre> |

4. Propagate YP Maps

"Propagate YP maps" means to copy a map from one server to another. Initially, *ypinit(1M)* copies the maps when you create slave servers.

After the slave servers are initialized, you will use *ypxfr(1M)* to transfer updated maps from the master server to the slaves. You can run *ypxfr(1M)* three ways:

- periodically from *cron(1M)* on each slave server,
- periodically by executing *yppush(1M)* on the master server, or
- interactively executing *ypxfr(1M)* on a slave server.

crontab(1)

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(1)* entries to periodically run *ypxfr(1M)* at a rate appropriate for each map in the YP database. The *ypxfr(1M)* 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(1)* 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(1M)*. 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 YP domain. Alter the exact time of execution from one server to another to prevent this process from slowing down the master.

EXAMPLE: *crontab(1)* 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(1M)* from within your *crontab(1)* file.

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

yppush(1M)

Execute *yppush(1M)* only on the master server to copy a map to each server in the YP domain (retrieved from the *ybservers* map).

1. The *yppush(1M)* command sends a “transfer map” request to each of the servers.
2. In turn, *ybserv(1M)* on each server executes *ypxfr -C*.
3. The *ybserv(1M)* daemon then passes *ypxfr(1M)* the information needed to identify and transfer the map.

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

ypxfr(1M)

Execute *ypxfr(1M)* interactively 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 server consistent with the other servers.

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

If you want the map transferred from a server other than the master, specify it using the *-h* option with *ypxfr(1M)*.

EXAMPLE:

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

Verify YP

To verify a client is bound to a server, login to that client and execute *ypwhich(1)*.

- 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.

YP domain *domain_name* not bound.

If you try *ypwhich(1)* several times and continue to receive the *not bound* response, the node is unable to locate a server for that YP domain on the network. Review your YP configuration process. If you did not make errors, refer to the “Troubleshooting” chapter.

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

Note You have now completed configuring YP. If you are configuring YP 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, return to the section, “7. Execute */etc/netnfsrc*” in the “NFS Configuration and Maintenance” chapter.

YP Maintenance

To keep YP 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 YP
- Modify YP Maps
- Add New YP Servers
- Add New Users to a Node
- Make a Different Node the YP Master
- Change YP Password
- Log Files
- Create Non-standard YP Maps

Disable YP

You must be super user to disable YP. If you choose to disable the YP service, use the following steps:

1. Set the YP domain name to null (no spaces).

```
domainname
```

2. If the YP service is currently running, kill the *ybind(1M)* and *ypserv(1M)* processes.

3. Edit */etc/netnfsrc* to change the YP values.

a. Change the *YP_MASTER_SERVER*, *YP_SLAVE_SERVER*, and *YP_CLIENT* values to zero.

```
YP_MASTER_SERVER=0
YP_SLAVE_SERVER=0
YP_CLIENT=0
```

b. Remove the *YPDOMAIN* variable if one exists.

```
YPDOMAIN=
```

4. If the above YP domain is specified in */etc/netgroup*, remove the YP domain name throughout this file.

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

6. Reboot the system.

Modify YP Maps

You must be super-user to modify YP maps.

Note 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 YP maps, like */etc/hosts*, by first editing the ASCII file and then running *ypmake(1M)*. Refer to the following “Manual Modifications to YP Maps” section if you are

- adding non-standard maps,
- editing maps for which no ASCII file exists, or
- changing the set of servers after the system is running.

Whether using *ypmake(1M)* 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 YP domain directory on the master server. (Refer also to *makedbm(1M)*).

Note **Never** modify a map directly; **always** use *makedbm(1M)* to create the map.

Manual Modifications to YP 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)

1. Move to the directory in which the maps you want to modify exist.

```
cd /user/etc/yp/YP_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(IM)* 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(IM)* output which you can then direct as input to *makedbm(IM)*. 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
```

EXAMPLE: Suppose you want to create a non-standard YP map. You want it to consist of key-value pairs in which the keys are strings like *al*, *bl*, *cl*, and *dl*, and the values are *ar*, *br*, *cr*, and *dr*. After creating the map, you notice it is missing *dl* 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:
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 YP domain subdirectory where the map is located
- The YP 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 *dl* 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 dl 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 YP domain subdirectory in which you will create the map

First, create the YP map from the keyboard by entering input on the master server as follows.

```
cd /usr/etc/yp
./makedbm - reports_domain/wes_map
al ar
bl br
cl cr
CTRL-D
```

To modify the map, use *makedbm(1M)* 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 *dl* 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 New YP Servers

You must be super-user to add new YP 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 also. (Refer to *ypmake(1M)*.)

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(1M)*.

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

2. Add the host's name to the *ypservers* map in the YP domain as follows.

```
cd /usr/etc/yp  
./makedbm -u home_domain/ypservers;\  
echo new_ypslave_name | ./makedbm - home_domain/ypservers  
yppush ypservers
```

3. On the new slave server, complete the steps in the “3. Create a YP Slave Server” section.

Add New Users to a Node

You must be super-user to add new users to a node.

Refer to the *HP-UX System Administrator's Manual/HP 9000 Series 800* or *HP-UX System Administration Tasks/HP 9000 Series 300* 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, and (3) defining the new user's environment.

Remember to update the YP *passwd* and *group* databases by running */usr/etc/yp/ypmake*. If you are using an alternate file to build the YP *password* databases, use its full path name instead of */etc/passwd*.

```
/usr/etc/yp/ypmake group passwd PWFIL = alternate_passwd_file
```

Make a Different Node the YP Master

You must be super-user to change the YP 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 pass word 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 YP Master Server" section.
6. To prevent starting *yppasswdd(1M)* on the old master server, edit its */etc/netnfsrc* file to change the *YP_MASTER_SERVER* value to zero.

```
YP_MASTER_SERVER=0
```

7. If the old master server is to be a slave server, complete the steps in the "3. Create a YP Slave Server" section and the steps in the "2. Create a YP 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(1M)* on each slave server using the new master server's host name.

```
ypinit -s new_master_hostname
```

Create or Change YP Password

The YP password is the password for a user's login ID that exists in the YP *passwd* map. The YP password is used as the user password, but is administered through YP. Note, you do not have to have a YP password to access the YP database.

If you change your password with the *passwd(1)* 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 YP is in use, execute *yppasswd(1)*.

YP Password Guidelines

The following list provides the requirements for creating and changing YP passwords.

- Only the owner or super-user can change a YP password. The super-user must know the current YP password to change another user's YP password.
- Only the first eight characters of the YP password are significant; the rest are truncated.

- A YP password must contain at least five characters if it includes a combination of either
 - uppercase and lowercase lettersor
 - alpha-numeric characters.
- A YP password must contain at least four characters if it includes a combination of uppercase letters, lowercase letters, and numeric characters.
- A YP password must contain at least six characters if it includes only monospace letters.
- You can change a YP password in the YP *passwd* map using *yppasswd(1)* only if *rpc.yppasswdd* is running on the master server. (See *yppasswdd(1M)*.)

YP Password

Refer to the following steps to create or change your YP password in the YP *passwd* map.

1. Execute the *yppasswd* command.

```
yppasswd user_login_name
```

2. The system prompts you for the old YP password even if one does not exist. If it does exist, enter the old YP password; otherwise, press RETURN.

Note, the YP password may be different from the one on your local node.

3. The system prompts you for the new YP password twice to ensure you enter the correct response. Enter your new YP 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(1M)*, *ypserv(1M)*, and *yppasswdd(1M)* so that diagnostic and error messages are written to log files.

```
/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(1M)* daemon writes its messages directly to the system console, */dev/console*.
- The *ypserv(1M)* daemon writes its messages to the */usr/etc/yp/ypserv.log* file if it exists when *ypserv(1M)* is started.
- The *yppasswdd(1M)* daemon provides no messages.

The *ypxfr(1M)* 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(1M)* is not being run directly by someone at a terminal.

EXAMPLE: Logging occurs if the log file exists and *cron(1M)* is running *ypxfr(1M)* directly, using a *crontab(1)* entry like the following one.

```
30 * * * * /usr/etc/yp/ypxfr yp_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(1)* 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 YP Maps

You must be super-user to create and propagate non-standard YP 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.

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 YP 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)*.)
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(1M)* 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(1)* entries or adding new entries to one of the *ypxfr(1M)* shell scripts described in the "Propagate YP Maps" section.

EXAMPLE:

This example spans several pages. Refer to the bold face headings for different sections of the example.

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 YP 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 login 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
```

The above *ypmatch* command would provide 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(1M)` 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 YP servers.

Maintain a Current Access Map on Each Slave Server

1. Execute the newly modified *ypmake(1M)* 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(1M)* script to periodically copy the *access* map from the master server.

```
# ypxfr_1perday - Perform daily YP 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 YP 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.

Note, the order of the *ypcat(1)* 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(1)* 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  betacatfish
mr_jad     axesys handel
```

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 Yellow Pages (YP) service is not mandatory for using VHE, but this chapter shows how to use VHE assuming YP is configured and used.

If you do not plan to use YP, 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 YP (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 YP master server using the information from step 3.
5. Edit the */etc/passwd* file on the YP 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 master YP 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.

Configuration

The following sections describe the configuration steps in detail.

Note You must be super-user 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 YP (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 YP 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 YP, refer to the “YP Configuration and Maintenance” chapter. VHE can use YP 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. YP provides this; if not using YP, 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*.”)

YP maintains single versions of the */etc/passwd* and */etc/vhe_list* files on the YP master server. From the YP 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 YP maps and propagated to the YP 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.

- a. Compare each */etc/newconfig* file listed below with its counterpart shown in the following list.

| File in <i>/etc/newconfig</i> directory | Counterpart in <i>/usr/etc/vhe</i> directory |
|---|--|
| <i>vhe_mounter</i> | <i>vhe_mounter</i> |
| <i>vhe_script</i> | <i>vhe_script</i> |

- b. If the files are the same, skip to the next section, “Determine File Systems and Mount Point Directories.”
- c. 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*.

OR

- 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 super-user, edit a file named *vhe_list* in the */etc* directory of the YP 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 YP, 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 YP 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
```

The following example shows */etc/passwd* file entries before and after the VHE configuration.

Note If you are not using YP, after updating the */etc/passwd* file, you must distribute the changes to all nodes in the VHE group.

Example

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/jbr1:/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/jbr1* to */vhe/C/usr/mail/jbr1* (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 YP that are part of the same YP domain), execute the following command on the YP master server.

```
/usr/etc/yp/ypmake
```

This builds the YP maps and propagates the maps to the YP 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 super-user 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 | | | | |
|----------|---------------------|---------------------|---------------------|---------------------|
| A | /vhe/A symlink=/ | /vhe/B Directory | vhe/C Directory | /vhe/D Directory |
| B | /vhe/A Directory | /vhe/B symlink=/ | /vhe/C Directory | /vhe/D Directory |
| C | /vhe/A Directory | /vhe/B Directory | /vhe/C symlink=/ | vhe/D Directory |
| D | /vhe/A Directory | /vhe/B Directory | /vhe/C Directory | /vhe/D 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, return to the “7. Execute `/etc/netnfsrc`” section in the “NFS Configuration and Maintenance” chapter.

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 YP 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 YP servers, execute the following command on the YP 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 super-user. 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*.

- 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
```

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), or
- 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 preceded by **YP**.

Key Terms

| Term | Definition |
|------------------------------|---|
| Client | <ul style="list-style-type: none"> ■ A node that requests data or services from other nodes (servers). ■ A process that requests other processes to perform operations. <p>Note: An NFS client can also be configured as any combination of an NFS server, YP client, or YP server. (A YP server must also be configured as a YP client.)</p> |
| Cluster | One or more workstations linked together with a local area network (LAN), but consisting of only one file system. |
| Cnode | Any node operating in an HP-UX cluster environment, including diskless nodes and the root server. |
| Daemon | Background programs that are always running, waiting for a request to perform a task. |
| Diskless Cnode | A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a file system directly attached. |
| 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. |
| Host | A node that has primary functions other than switching data for the network. |
| Heterogeneous Cluster | A diskless cluster with more than one type of computer attached. |
| Homogenous Cluster | A diskless cluster composed of nodes of only one type of computer architecture (e.g., HP 9000 Series 300) |

| Term | Definition |
|---------------------------|--|
| Map (YP) | <p>A file consisting of logical records; a search key and related value form each record. YP clients can request the value associated with any key within a map.</p> <p>YP map is synonymous with YP database.</p> |
| Master Server (YP) | <p>The node on which one or more YP maps are constructed from ASCII files. These maps are then copied to the YP slave servers for the YP clients to access.</p> |
| Mount | <p>To obtain access to a remote or local file system or directory (import).</p> |
| Mount Point | <p>The name of the directory on which a file system is mounted.</p> |
| Netgroup | <p>A network-wide group of nodes and users defined in <i>/etc/netgroup</i>.</p> |
| Node | <p>A computer system that is attached to or is part of a computer network.</p> |
| Root Server | <p>The only node in an HP-UX cluster that has file systems physically attached to it.</p> |
| Server | <ul style="list-style-type: none"> ■ A node that provides data or services to other nodes(clients) on the network. ■ A process that performs operations as requested by other processes. <p>Note: An NFS server can also be configured as any combination of an NFS client, YP client, or YP server. (A YP server must also be configured as a YP client.)</p> |
| Slave Server (YP) | <p>A node that copies YP maps from the YP master server and then provides YP clients access to these maps.</p> |
| Soft Mount | <p>An optional mount that causes access to remote file systems to abort requests after one NFS attempt.</p> |

| Term | Definition |
|--------------------------|--|
| Yellow Pages (YP) | <p>An optional network service composed of databases (maps) and processes that provide YP clients access to the maps. The YP service enables you to administer these databases from one node.</p> <p>YP may or may not be active; check with your system administrator.</p> |
| YP Client | <ul style="list-style-type: none"> ■ A node that requests data or services from YP servers. ■ A YP process that requests other YP processes to perform operations. <p>Note: A YP client can also be configured as any combination of a YP server, NFS client, or NFS server. (A YP server must also be configured as a YP client.)</p> |
| YP Database | See “Map (YP).” |
| YP Domain | A logical grouping of YP maps (databases) stored in one location. YP domains are specific to the YP network service and are not associated with other network domains. |
| YP Map | See “Map (YP).” |
| YP Password | <p>The password for a user’s login ID that exists in the YP <i>passwd</i> map. The YP password is the same one as the user password, but is administered through the YP.</p> <p>You do not have to have a password to access the YP databases.</p> |

| Term | Definition |
|-----------|--|
| YP Server | <ul style="list-style-type: none"><li data-bbox="491 183 1163 244">■ A node that provides data (maps) or services to other nodes (YP clients) on the network using YP.<li data-bbox="491 257 1049 318">■ A YP process that performs operations as requested by other YP processes. <p data-bbox="534 353 1120 447">Note: A YP server must also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.</p> |

Troubleshooting References

Troubleshooting the NFS Services primarily concerns the areas: power up and connectivity, NFS Services, YP Services, VHE, and REX. This chapter only addresses NFS, YP, 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.

- *HP 98643A LAN/300 Link LANIC Installation Manual*
- *HP Repeater Installation Manual* (only if you are using a HP 92223A repeater)
- *HP-UX Installation Manual*
- *HP-UX Reference manuals*
- *HP-UX System Administrator's Manual/HP 9000 Series 800*
- *HP-UX System Administration Tasks/HP 9000 Series 300*
- *LAN Cable and Accessories Installation Manual*
- *Installing and Administering NS Services*
- *Installing and Administering ARPA Services*

Troubleshooting Sections

Refer to the “Troubleshoot NFS” section or the *NFS Services Reference Pages* 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 “Troubleshoot Yellow Pages” section or the *NFS Services Reference Pages* if you configured the system to use YP, but cannot access files serviced by it.

Refer to the “Troubleshoot 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 “Troubleshoot 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 and HP-UX 2.0 or later for the Series 800? For File Locking and REX, your host must be running HP-UX 6.5 or later for the Series 300 or HP-UX 7.0 or later for the Series 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 or Series 800? If so, refer to the appropriate system documentation.

Hardware

The *Installing and Administering NS Services* and *Installing and Administering ARPA Services* 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 Ethernet¹ hardware with IEEE 802.3² 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?
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, Installing and Administering NS Services, and Installing and Administering ARPA Services* documentation.
3. Have you performed the corrective action supplied with the error message you received? Consult the appropriate entry in the network reference pages.

(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.

4. If using gateways, do both hosts have routing information to each other? Refer to the *ARPA/Berkeley Services Reference Pages*, *route(1M)* section.
5. **If operating in an HP-UX cluster environment** and trying to mount an NFS file system, ensure you are using the root server's host name as the node specified in the *mount(1M)* command. Note, you can troubleshoot NFS specific problems from the root 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 NS Services* and *Installing and Administering ARPA Services* documentation.)

NFS and Yellow Pages

1. Is the client system trying to perform tasks as super-user 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 *HP-UX System Administrator's Manual/HP 9000 Series 800*, or the *HP-UX System Administration Tasks/HP 9000 Series 300* manual, and in the "NFS Configuration and Maintenance" and "YP 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 and HP-UX release 7.0 for the Series 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 and HP-UX release 7.0 for the Series 800, HP-UX did **not** support remote access to device files.
5. Does the *inetd* security file (*/usr/adm/inetd.sec*) contain 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(1M)* 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 YP 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 *rex*d? The server must have an entry in */etc/inetd.conf* in order to run *rex*d (see *rex*d(1M)).
2. Was network communication established between the REX client and the REX server using procedures outlined in the *HP-UX System Administrator's Manual/HP 9000 Series 800*, or the *HP-UX System Administration Tasks/HP 9000 Series 300* manual, and in the "NFS Configuration and Maintenance," "YP 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 *rex*d service?
4. Does the user have a user account on both the REX client and the REX server with matching UIDs?

5. Was `rex` 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 YP 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 *NFS Services Reference Pages* 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

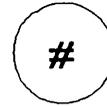
If you do not solve the problem after working through the previous troubleshooting steps and following flowcharts, call your HP support representative for assistance. Provide as much information about the problem as possible, including information from your network map and the following items.

- 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.

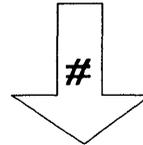
Flowchart Format

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.

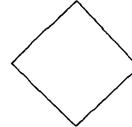
Start of Flowchart #



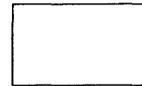
Go to and enter
specified Flowchart#



Make a decision



Perform an action



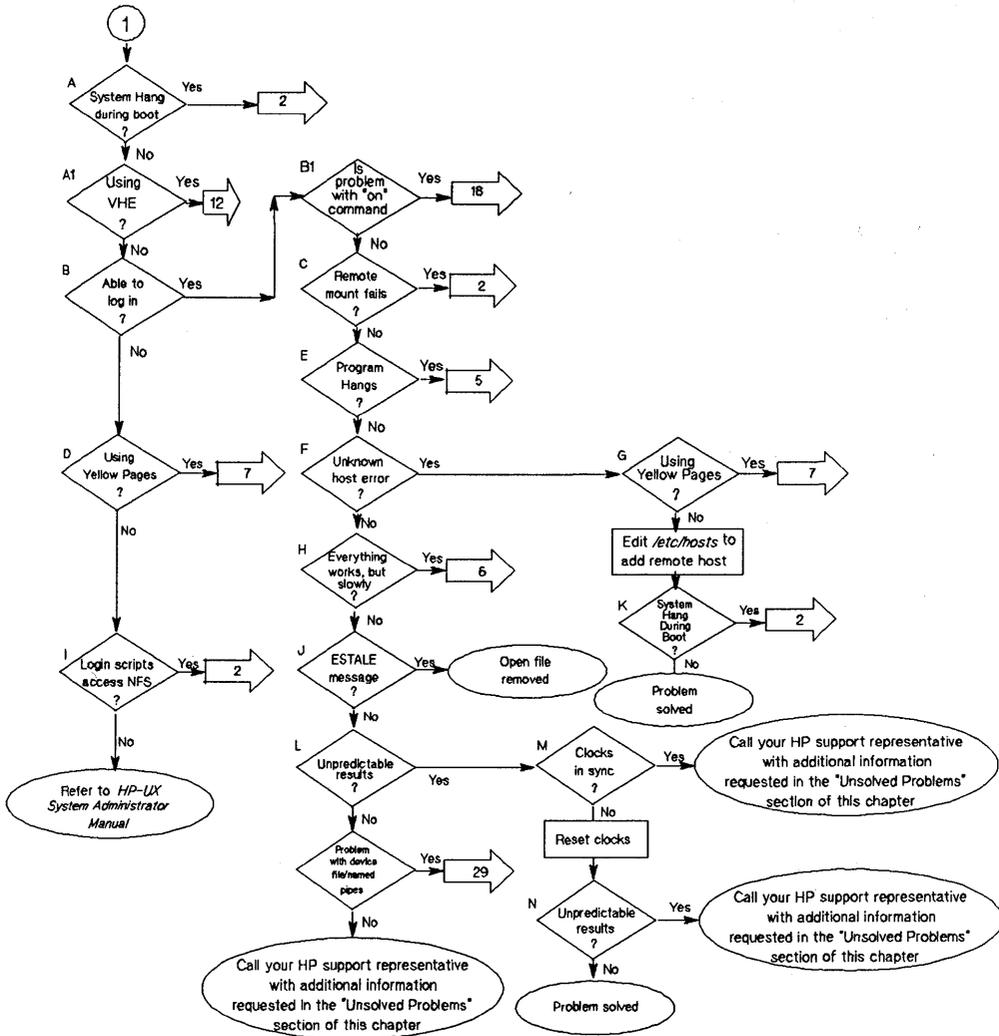
Exit Flowchart



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(1M)*).

Troubleshoot NFS



Flowchart 1: Initial Steps to Narrowing the Problem

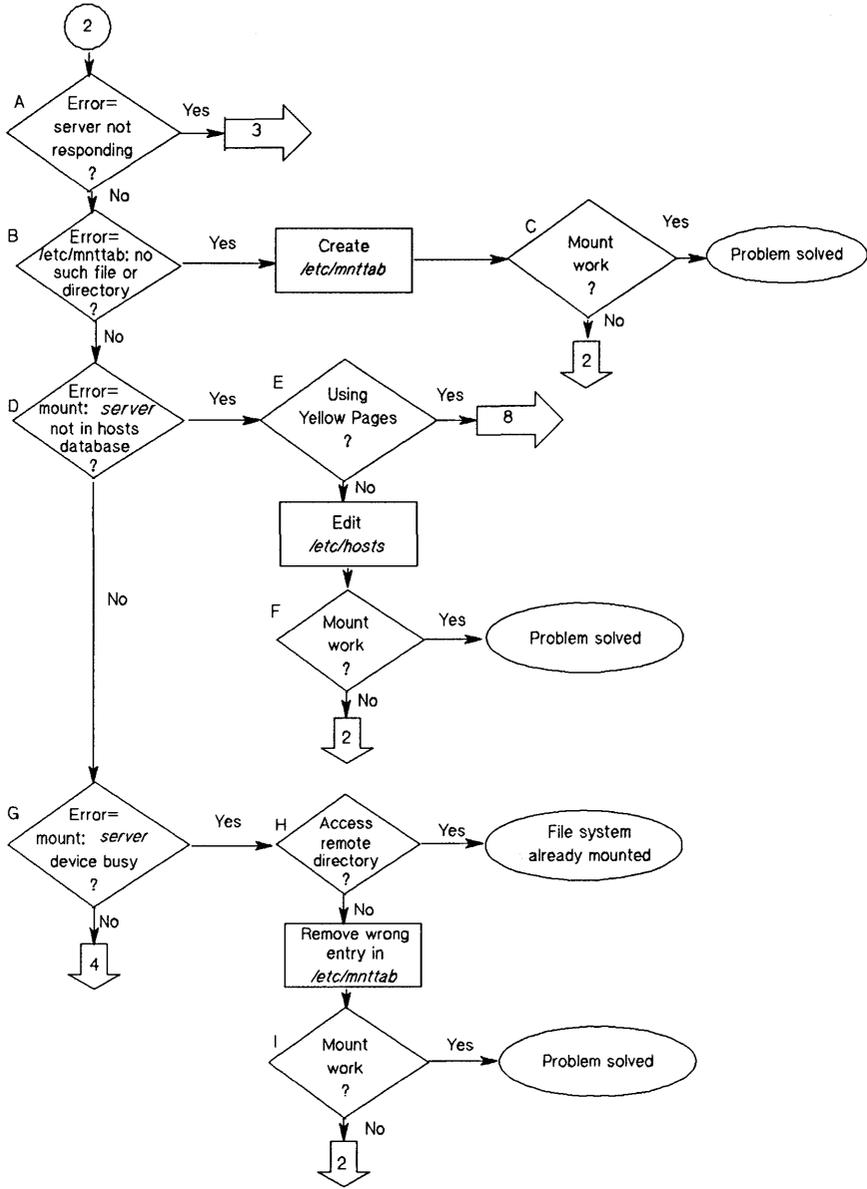
Initial Steps to Narrowing the Problem (Flowchart 1)

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 |
|--|-------------------|---|
| <p>A. Does the system hang during boot when mounting remote files?</p> <p>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.</p> | See Flowchart 2. | See A1. |
| <p>A1. Are you using VHE?</p> | See Flowchart 12. | See B. |
| <p>B. Are you able to login?</p> | See C. | You will receive error messages or the system will fail to respond if you cannot log in to it. See D. |
| <p>B1. Is the problem experienced while using the "on" command?</p> | See Flowchart 18. | See C. |
| <p>C. When trying to mount a remote file system, do error messages indicate the attempt failed?</p> | See Flowchart 2. | See E. |
| <p>D. Are you using YP?</p> | See Flowchart 7. | See I. |

| Question | Yes: Action | No: Action |
|---|--|---|
| E. Do programs performing remote file accesses hang? | See Flowchart 5. | See F. |
| F. Does the system report unknown host errors during execution of commands or programs? | See G. | See H. |
| G. Are you using YP? | See Flowchart 7. | Edit <i>/etc/hosts</i> 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 <i>HP-UX System Administrator's Manual/HP 9000 Series 800</i> or the <i>HP-UX System Administration Tasks/HP 9000 Series 300</i> 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. |

| Question | Yes: Action | No: Action |
|--|---|--|
| 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 <i>date(1)</i> command, and then see N. |
| 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

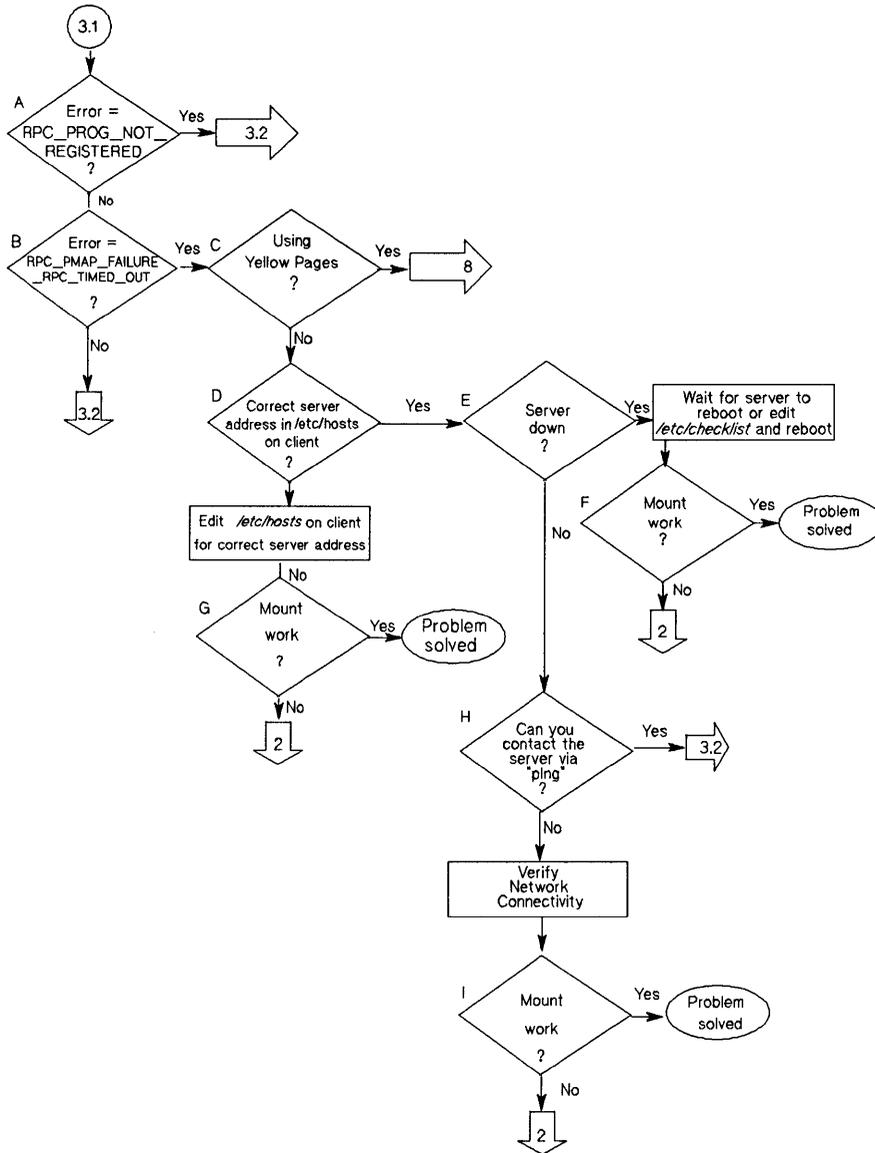
Mount Fails (Flowchart 2)

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(1M)* 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. | See B. |
| B. Does the following error message occur on the client? <i>/etc/mnttab</i> : no such file or directory | Create <i>/etc/mnttab</i> on the client, and then see C. The system uses <i>/etc/mnttab</i> to log all mounted file systems. Note: Generally, at boot time <i>/etc/rc</i> creates <i>/etc/mnttab</i> . | See D. |
| C. Can you mount the remote system? | Problem solved. | Restart Flowchart 2. |
| D. Does the following error message occur on the client? <i>mount: server</i> not in hosts database | See E. | See G. |
| E. Are you using YP? | See Flowchart 8. | Edit <i>/etc/hosts</i> on the client to include the desired remote host, and then see F. |
| F. Can you mount the remote system? | Problem solved. | Restart Flowchart 2. |
| G. Does the following error message occur on the client? <i>mount: server</i> device busy | See H. | See Flowchart 4. |

| Question | Yes: Action | No: Action |
|--|---|---|
| H. 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 <i>/etc/mnttab</i> for the remote file system you are trying to mount, and then see I . |
| I. Can you mount the remote system? | Problem solved. | Restart Flowchart 2. |



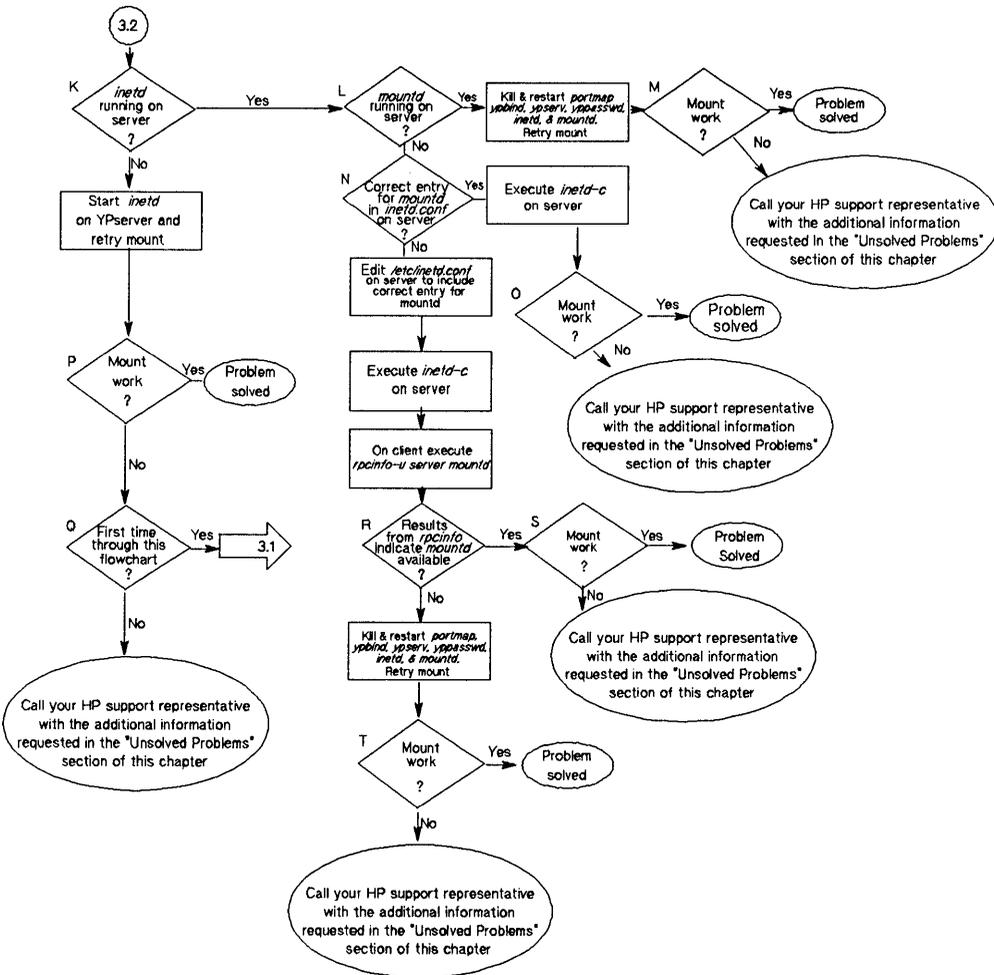
Flowchart 3.1: Server Not Responding

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 YP? | See Flowchart 8. | See D. |
| D. Is the server's address correct in the client's <i>/etc/hosts</i> ? | See E. | Edit the client's <i>/etc/hosts</i> to include the correct address for the server you are trying to mount. See G. |
| E. 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: <ul style="list-style-type: none"> ■ Do nothing on the system until the server reboots. ■ Edit the client's <i>/etc/checklist</i> to remove the NFS entry for that server; reboot the system. See F. | See H. |
| F. Can you mount the remote system? | Problem solved. | See Flowchart 2. |
| G. Can you mount the remote system? | Problem solved. | See Flowchart 2. |

| Question | Yes: Action | No: Action |
|---|---------------------------|---|
| <p>H. Can you contact the server using the ping diagnostic? Refer to the <i>Installing and Administering NS Services</i> and <i>Installing and Administering ARPA Services</i> manuals for <i>ping</i> diagnostic information</p> | <p>See Flowchart 3.2.</p> | <p>Refer to the <i>Installing and Administering LAN</i> manual to verify Link connectivity, and then see I.</p> |
| <p>I. Can you mount the remote system?</p> | <p>Problem solved.</p> | <p>See Flowchart 2.</p> |



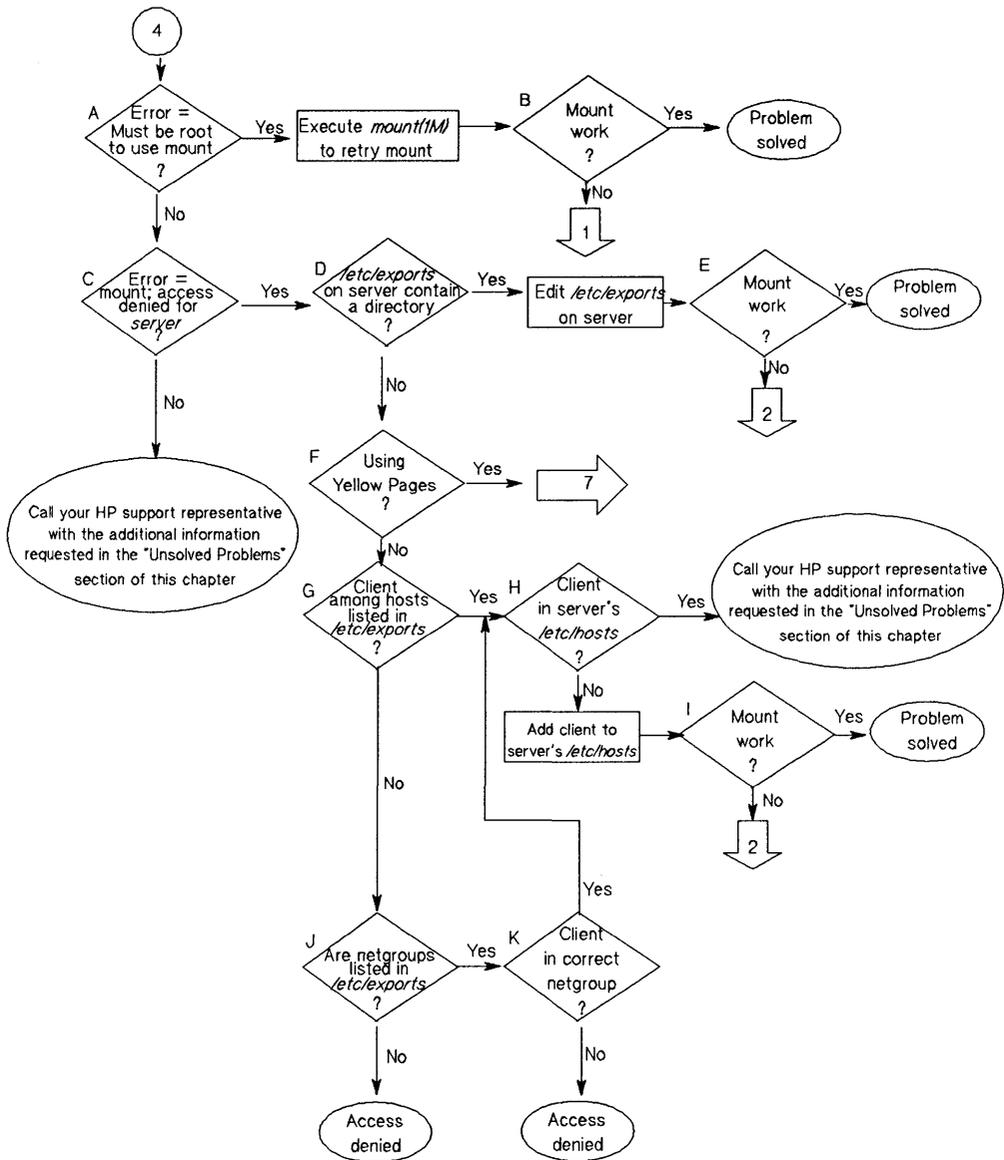
Flowchart 3.2: Server Not Responding

Server Not Responding (Flowchart 3.2)

| Question | Yes: Action | No: Action |
|--|--|---|
| K. Is <i>inetd(1M)</i> running on the server? | See L. | Start <i>/etc/inetd</i> on the server, retry the mount, and then see P. |
| L. Is <i>mountd(1M)</i> running on the server? | Kill and restart the following daemons on the server in the order specified. <ul style="list-style-type: none"> ■ <i>portmap(1M)</i> ■ <i>ybind(1M)</i> * ■ <i>ypserv(1M)</i> * ■ <i>yppasswdd(1M)</i> * ■ <i>inetd(1M)</i> ■ <i>mountd(1M)</i> * only if using YP 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. |
| N. Is the correct <i>mountd(1M)</i> entry in <i>inetd.conf</i> on the server? Ensure the entry is not commented out with a # (pound sign). | Execute <i>inetd -c</i> on the server, and then see O. | 1. Edit the server’s <i>/etc/inetd.conf</i> file to include the correct <i>mountd(1M)</i> entry. 2. Execute <i>inetd -c</i> on the server to read changes in <i>/etc/inetd.conf</i> . 3. Execute <i>rpcinfo -u</i> on the client. <code>rpcinfo -u server mountd</code> 4. See R. |

| Question | Yes: Action | No: Action |
|--|------------------------|---|
| 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. |
| R. Do the results from <i>rpcinfo -u</i> indicate a <i>mountd (1M)</i> process is available on the server? | See S. | <p>Kill and restart the following daemons on the server in the order specified.</p> <ul style="list-style-type: none"> ■ <i>portmap(1M)</i> ■ <i>ypbind(1M)</i> * ■ <i>ypserv(1M)</i> * ■ <i>yppasswdd(1M)</i> * ■ <i>inetd(1M)</i> ■ <i>mountd(1M)</i> <p>* only if using YP</p> <p>Retry the mount, and then see T.</p> |
| 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.. |

| Question | Yes: Action | No: Action |
|-------------------------------------|--------------------|---|
| 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. |

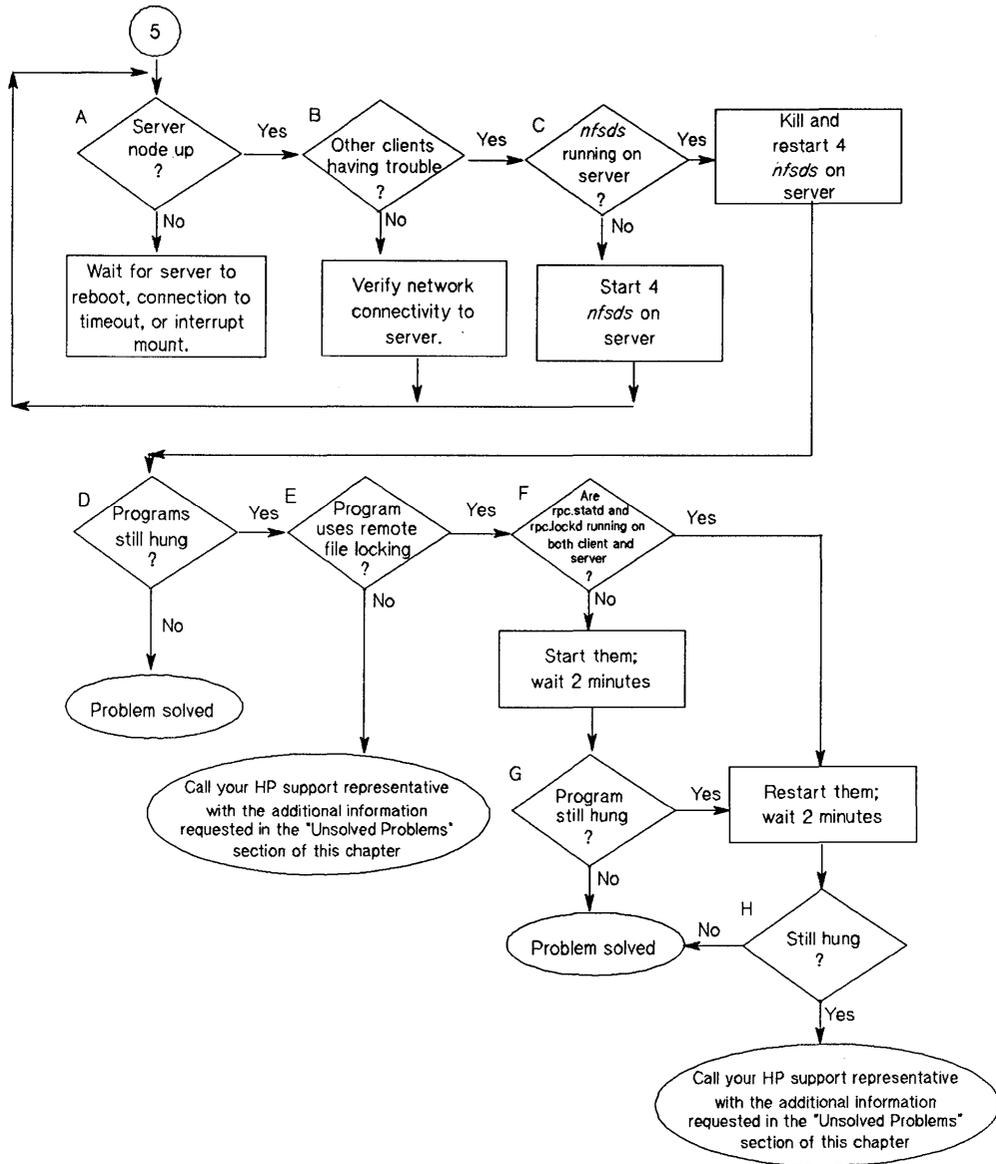


Flowchart 4: Restricted Access

Restricted Access (Flowchart 4)

| Question | Yes: Action | No: Action |
|--|---|---|
| <p>A. Does the following error message occur on the client?</p> <p>Must be root to use mount</p> | Login as super-user, execute <i>mount(1M)</i> , and then see B. | See C. |
| <p>B. Can you mount the remote system?</p> | Problem solved. | See Flowchart 1. |
| <p>C. Does the following error message occur on the client?</p> <p>mount: access denied for <i>server</i></p> | See D. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |
| <p>D. Does the server’s <i>/etc/exports</i> file list a directory rather than a file system?</p> | Edit the server’s <i>/etc/exports</i> to contain the file system rather than a directory, and then see E. | See F. |
| <p>E. Can you mount the remote system?</p> | Problem solved. | See Flowchart 2. |
| <p>F. Are you using YP?</p> | See Flowchart 7. | See G. |
| <p>G. If hosts are listed in <i>/etc/exports</i>, is the client among the hosts listed for the desired file system?</p> | See H. | See J. |
| <p>H. Is the client listed in the server’s <i>/etc/hosts</i>?</p> | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. | Add client to server’s <i>/etc/hosts</i>, and then see I. |
| <p>I. Can you mount the remote system?</p> | Problem solved. | See Flowchart 2. |

| Question | Yes: Action | No: Action |
|---|---------------|--|
| J. Are netgroups listed for this file system in server's <i>/etc/exports</i> ? | See K. | Access for this client is deliberately denied. |
| K. Is the client listed in the appropriate netgroup for this file system in <i>/etc/netgroup</i> ? | See H. | Access for this client is deliberately denied. |



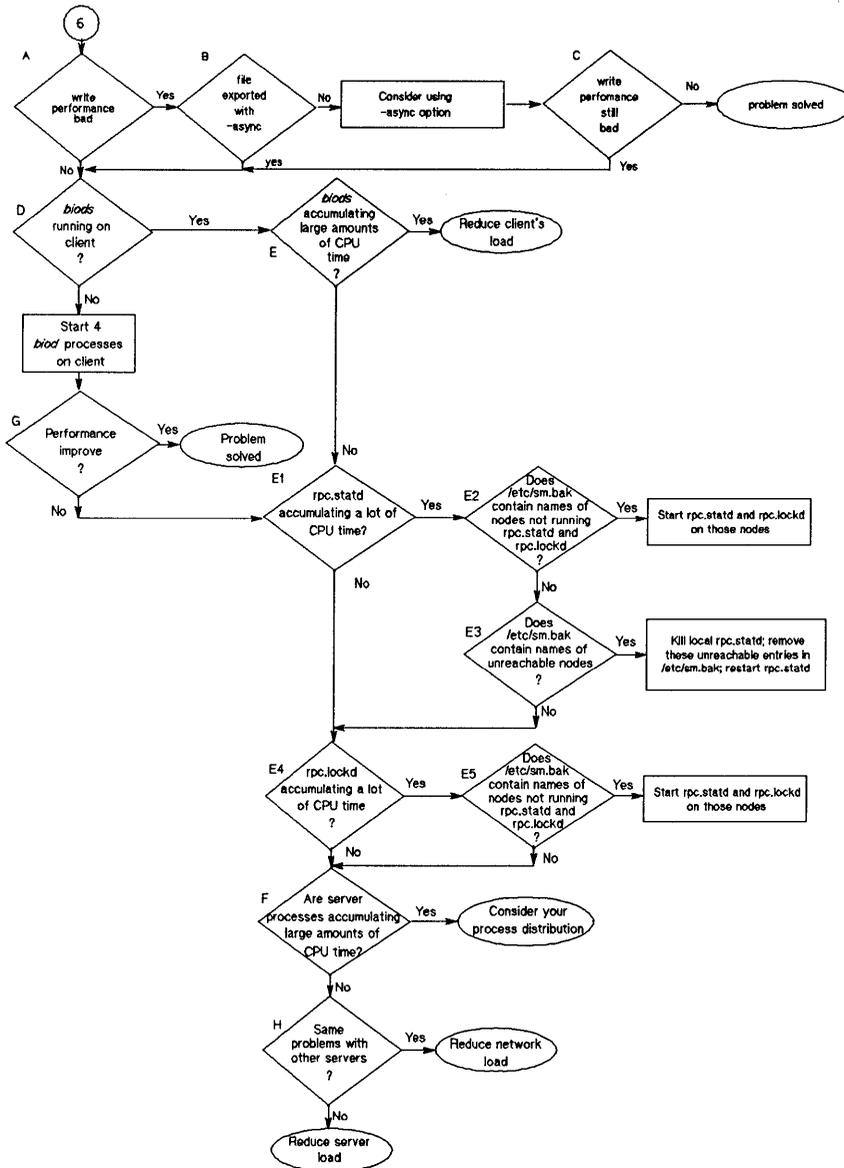
Flowchart 5: Programs Hang

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, either ■ wait for the server to reboot or ■ 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 <i>Installing and Administering LAN</i> manual. See A. |
| C. Are <i>nfsd(1M)</i> daemons running on the server? | Kill and restart four <i>nfsd(1M)</i> daemons on the server, and then see D. | Start four <i>nfsd(1M)</i> daemons on the server, and then see D. |
| 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. |
| F. Are <i>rpc.statd</i> and <i>rpc.lockd</i> running on both the client and the server? | Restart them and wait 2 minutes. See H. | Start them and wait 2 minutes. See G. |

| Question | Yes: Action | No: Action |
|---------------------------------|---|-----------------|
| 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. |



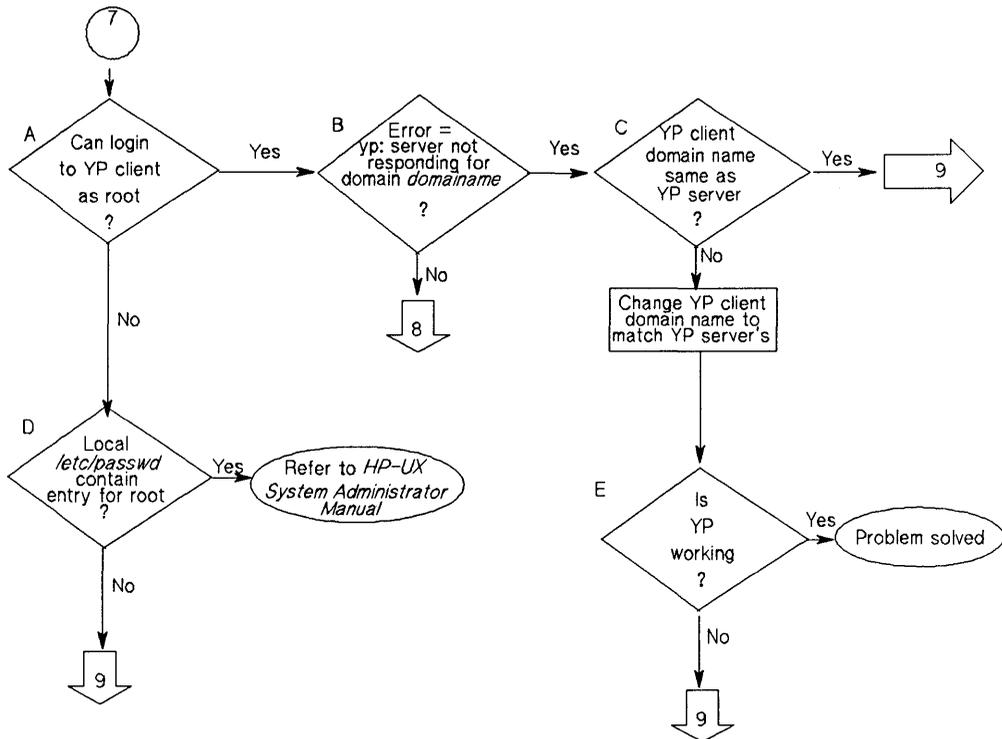
Flowchart 6: Performance Problems

Performance Problems (Flowchart 6)

| Question | Yes: Action | No: Action |
|---|---|--|
| <p>A. Are the <i>biod(1M)</i> daemons running on the client?</p> | <p>See B.</p> | <p>Start four <i>biod(1M)</i> processes on the client, and then see D.</p> |
| <p>B. Are the client <i>biod(1M)</i> daemons accumulating large amounts of CPU time?</p> <p>1. List the client processes using <i>ps</i>. 2. Copy a large file to the server system, and list the client <i>biod</i> processes again. 3. Compare the CPU time for the <i>biod(1M)</i> processes before and after the file copy.</p> | <p>Reduce the client's load to fewer NFS transactions by reducing the number of users or storing more files locally.</p> | <p>See B1.</p> |
| <p>B1. Is <i>rpc.statd</i> accumulating a lot of CPU time? (on the client?)</p> | <p>See B2.</p> | <p>See B4.</p> |
| <p>B2. Does <i>/etc/sm.bak</i> contain names of nodes not running <i>rpc.stad</i> and <i>rpc.lockd</i>?</p> | <p>Start <i>rpc.statd</i> and <i>rpc.lockd</i> on those nodes.</p> | <p>See B3.</p> |
| <p>B3. Does <i>/etc/sm.bak</i> contain names of unreachable nodes?</p> | <p>Kill local <i>rpc.statd</i>, remove these unreachable entries in <i>/etc/sm.bak</i>, and restart <i>rpc.statd</i>.</p> | <p>See B4.</p> |

| Question | Yes: Action | No: Action |
|--|---|--|
| B4. Is <i>rpc.lockd</i> accumulating a lot of CPU time (on the client)? | See B5. | See C. |
| B5. Does <i>/etc/sm.bak</i> contain names of nodes not running <i>rpc.statd</i> and <i>rpc.lockd</i> ? | Start <i>rpc.statd</i> and <i>rpc.lockd</i> on those nodes. | see C. |
| C. Are processes on the server accumulating large amounts of CPU time (especially <i>nfsd(1M)</i> , <i>inetd(1M)</i> , and <i>portmap(1M)</i> ?) | Consider whether you need to distribute your processing by adding additional systems. | See E. |
| D. Has performance improved? | Problem solved. | See B1. |
| E. Are the same performance problems evident with other servers? | Reduce the network load. | Reduce the server's load by adding more servers. |

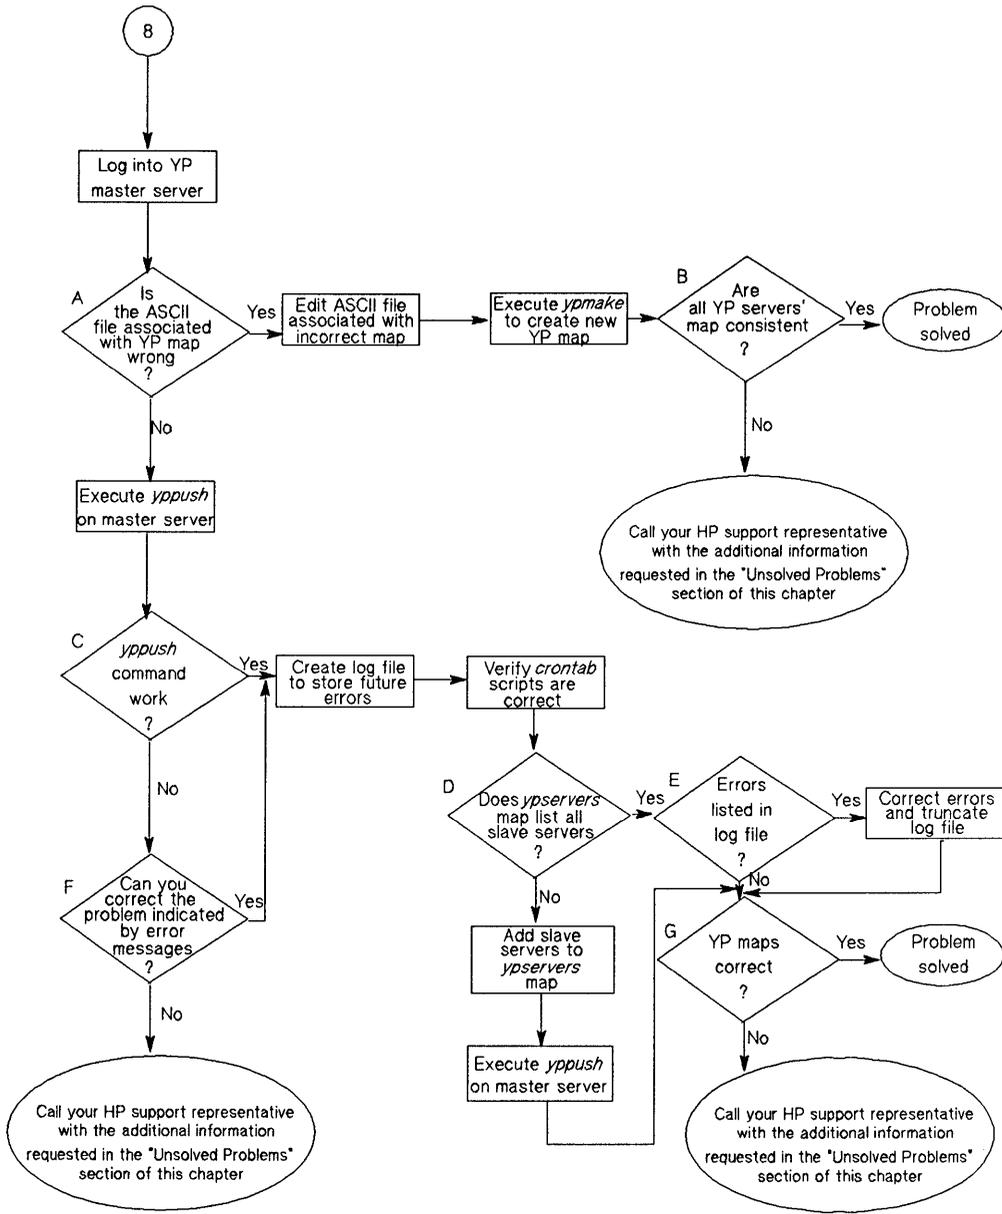
Troubleshoot Yellow Pages



Flowchart 7: Initial Steps to Troubleshooting YP

Initial Steps to Troubleshooting YP (Flowchart 7)

| Question | Yes: Action | No: Action |
|--|--|--|
| A. Can you login as root on the YP client? | See B. | See D. |
| B. Does the following error message occur on the console or in the <i>ybind</i> log file? yp: server not responding for domain <i>domain_name</i> | See C. | See Flowchart 8. |
| C. Is the YP client's YP domain name the same as the YP server's? | See Flowchart 9. | Change the YP client's YP domain name to be the same as the YP server's, and then see E. domainname <i>domain_name</i> |
| D. Does the local <i>/etc/passwd</i> file contain an entry for root? | The problem is not associated with YP or NFS. Refer to the <i>HP-UX Administrator's Manual</i> . | You cannot login to the YP client until YP is functioning unless you have an entry for a user in the local <i>/etc/passwd</i> file. See Flowchart 9. |
| E. Is YP working? If you can access the YP server's maps using <i>ypcat</i> or <i>ypmatch</i> , YP is probably functioning correctly. | Problem solved. | See Flowchart 9. |



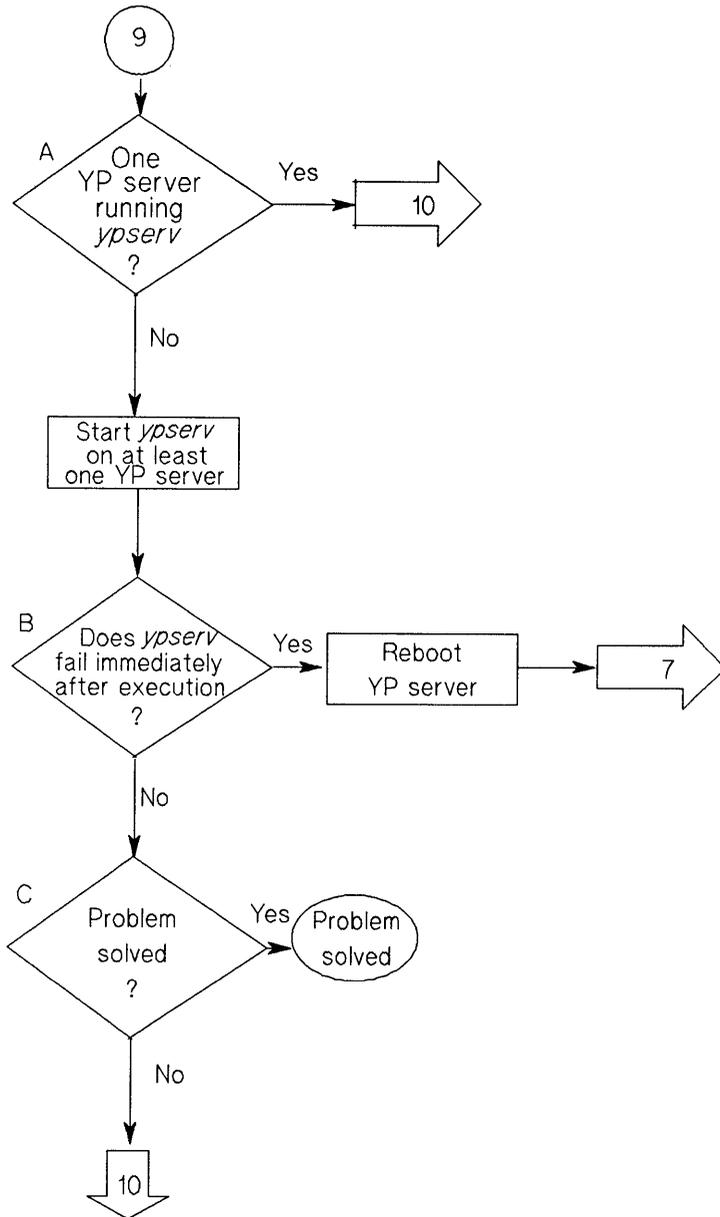
Flowchart 8: Incorrect YP Maps

Incorrect YP Maps (Flowchart 8)

Login to the YP master server as root before starting Flowchart 8.

| Question | Yes: Action | No: Action |
|---|---|--|
| <p>A. On the YP master server, does the ASCII file associated with the YP map need to be updated (e.g., update <i>/etc/hosts</i>)?</p> | <ol style="list-style-type: none"> 1. Edit the ASCII file associated with the incorrect YP map. 2. Execute <i>yppush(1M)</i> to create and distribute a new map to the YP slave servers. 3. See B. | <p>Execute <i>yppush(1M)</i> on the YP master server, and then see C.</p> <p><i>yppush map_name</i></p> |
| <p>B. Are all YP server's maps consistent? You can determine this by executing <i>yppoll</i> and then comparing order numbers.</p> | <p>Problem solved.</p> | <p>Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.</p> |
| <p>C. Does <i>yppush(1M)</i> work correctly? If you do not receive error messages associated with the command, it probably executed successfully.</p> | <ol style="list-style-type: none"> 1. Create the log file <i>/usr/etc/yp/ypxfr.log</i> to trap future errors associated with <i>yppush(1M)</i> on each YP slave server. 2. Verify that <i>crontab(1M)</i> scripts (on each slave server) copying the maps are correct. 3. See D. | <p>See F.</p> |

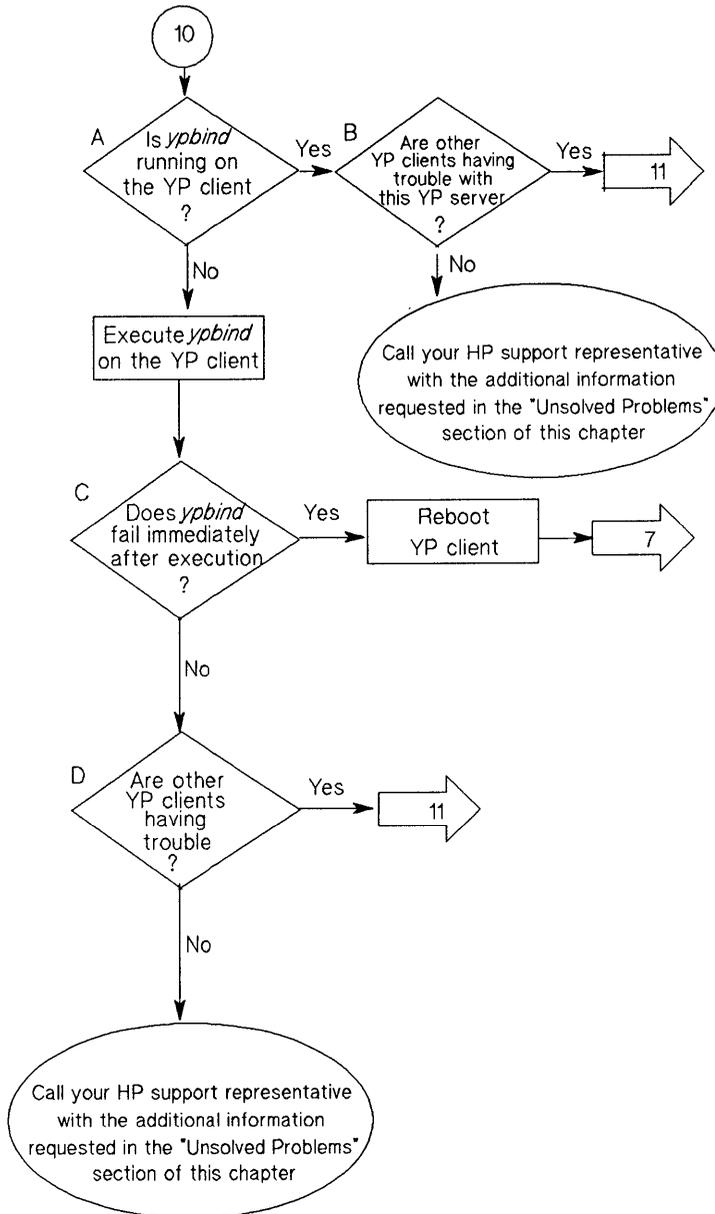
| Question | Yes: Action | No: Action |
|--|---|---|
| <p>D. Does the <i>ybservers</i> map list all YP slave servers? <code>yycat -k ybservers</code></p> | See E. | <p>1. Add any missing YP slave server to the <i>ybservers</i> map.</p> <p>2. Execute <i>yypush(1M)</i> on the YP master server to update all YP slave servers.</p> <p>3. See G.</p> |
| <p>E. Does <i>/usr/etc/yp/ypxfr.log</i> on the slave server list errors?</p> | Correct the errors, truncate the log file, and then see G. | See G. |
| <p>F. Can you correct the problem indicated by the error message?</p> | <p>1. Create the log file <i>/usr/etc/yp/ypxfr.log</i> to trap future errors associated with <i>yypush(1M)</i> on each YP slave server.</p> <p>2. Verify that <i>crontab(1M)</i> scripts (on each slave server) distributing the maps are correct.</p> <p>3. See D.</p> | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |
| <p>G. Are the YP maps correct?</p> | Problem solved. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |



Flowchart 9: ypserv(1M) Problems

ypserv(1M) Problems (Flowchart 9)

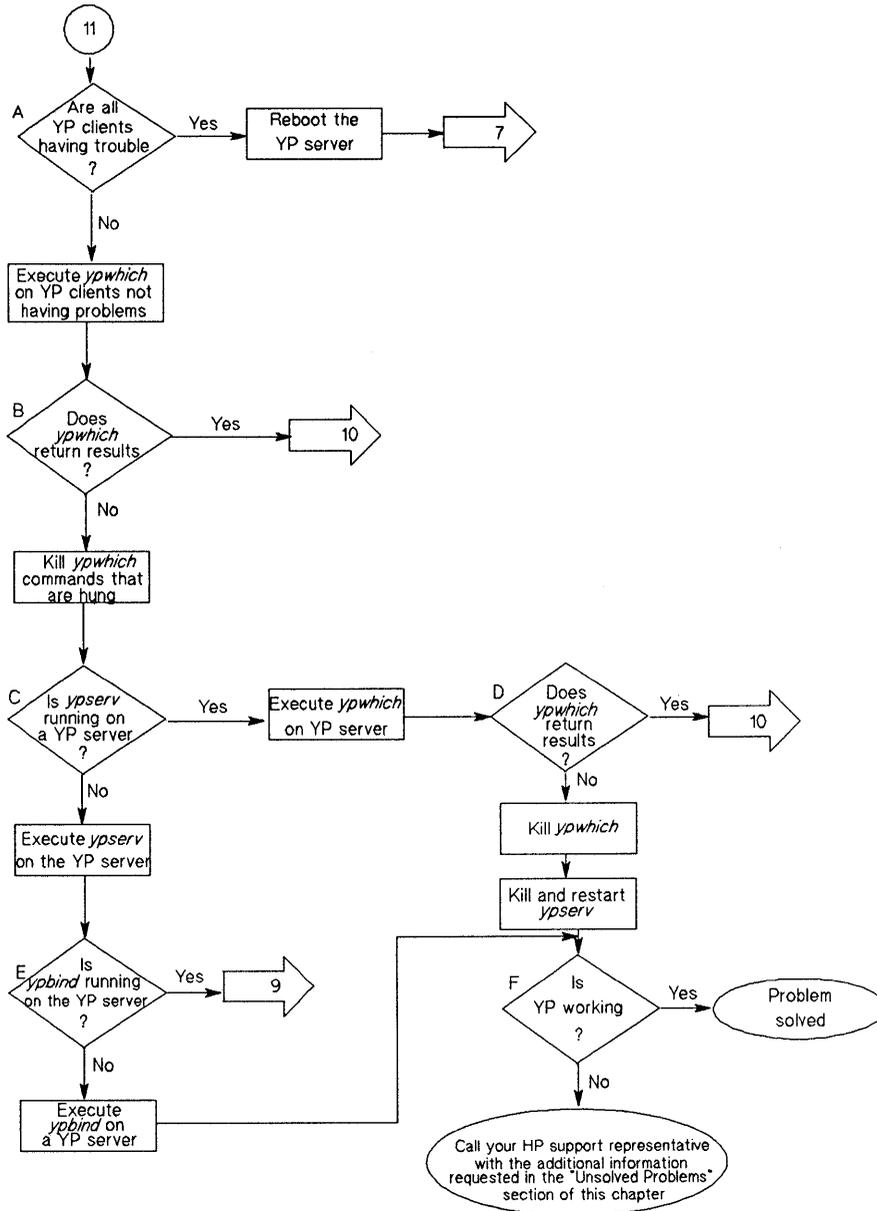
| Question | Yes: Action | No: Action |
|---|---|---|
| A. Is at least one YP server in the YP domain running <i>ypserv(1M)</i> ? | See Flowchart 10. | Start <i>ypserv(1M)</i> on at least one YP server in the YP domain, and then see B. |
| B. Does <i>ypserv(1M)</i> fail immediately after starting it? | Reboot the YP server, and then see flowchart 7. | See C. |
| C. Is the problem solved? | Problem solved. | See Flowchart 10. |



Flowchart 10: ypbind(1M) Problems

ypbind(1M) Problems (Flowchart 10)

| Question | Yes: Action | No: Action |
|--|---|---|
| A. Is <i>ypbind(1M)</i> running on the YP client? | See B. | Execute <i>ypbind(1M)</i> on the YP client, and then see C. |
| B. Are other YP clients having trouble with this YP server? | See Flowchart 11. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |
| C. Does <i>ypbind(1M)</i> crash immediately after starting it? | Reboot the YP client, and then see Flowchart 7. | See D. |
| D. Are other YP clients having trouble with this YP server? | See Flowchart 11. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |

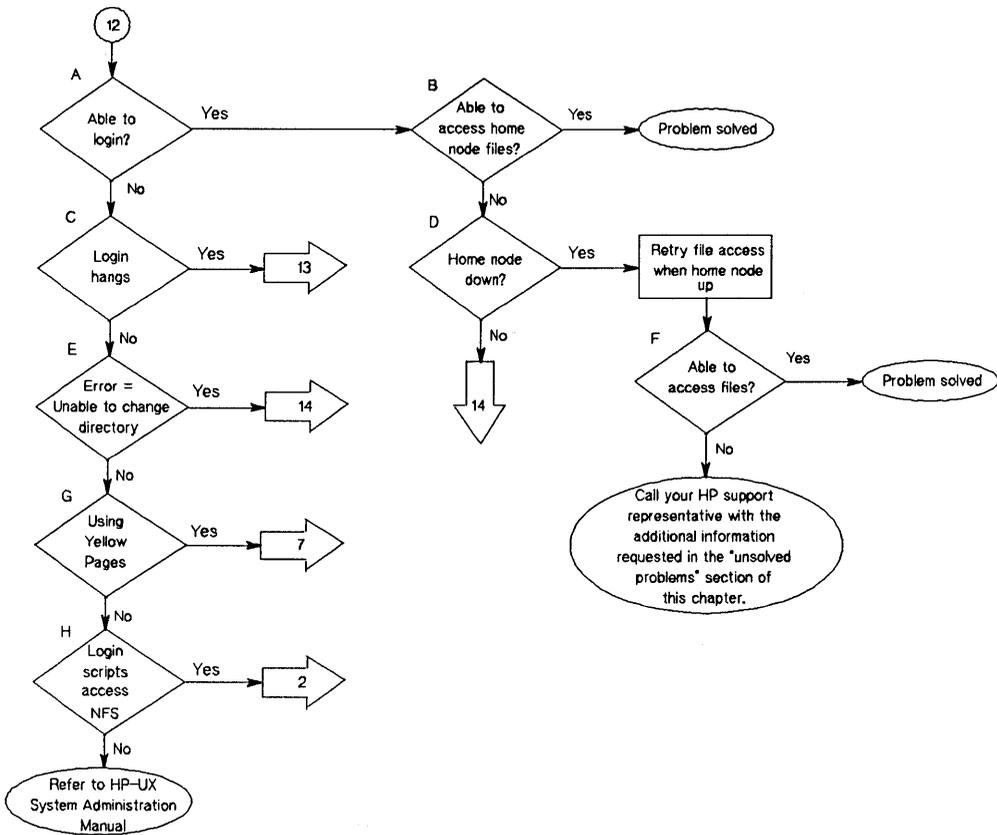


Flowchart 11: Multiple YP Client Problems

Multiple YP Client Problems (Flowchart 11)

| Question | Yes: Action | No: Action |
|--|---|--|
| A. Are all YP clients having trouble with this YP server? | Reboot the YP server, and then see Flowchart 7. | Execute <i>ypwhich(1)</i> on the YP client nodes not having problems, and then see B. |
| B. Does the <i>ypwhich(1)</i> command return results on the YP client? | See Flowchart 10. | Kill <i>ypwhich(1)</i> commands that are hung on YP clients, and then see C. |
| C. Is <i>ypserv(1M)</i> running on the YP server? | Execute <i>ypwhich(1)</i> on the YP server, and then see D. | Execute <i>ypserv(1M)</i> on the YP server, and then see E. |
| D. Does <i>ypwhich(1)</i> return results on the YP server? | See Flowchart 10. | <ol style="list-style-type: none"> 1. Kill <i>ypwhich(1)</i> on the YP server. 2. Kill and restart <i>ypserv(1M)</i>. 3. See F. |
| E. Is <i>ypbind(1M)</i> running on the YP server? | See Flowchart 9. | Execute <i>ypbind(1M)</i> on the YP server, and then see F. |
| F. Is YP functioning correctly on all YP clients? | Problem solved. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |

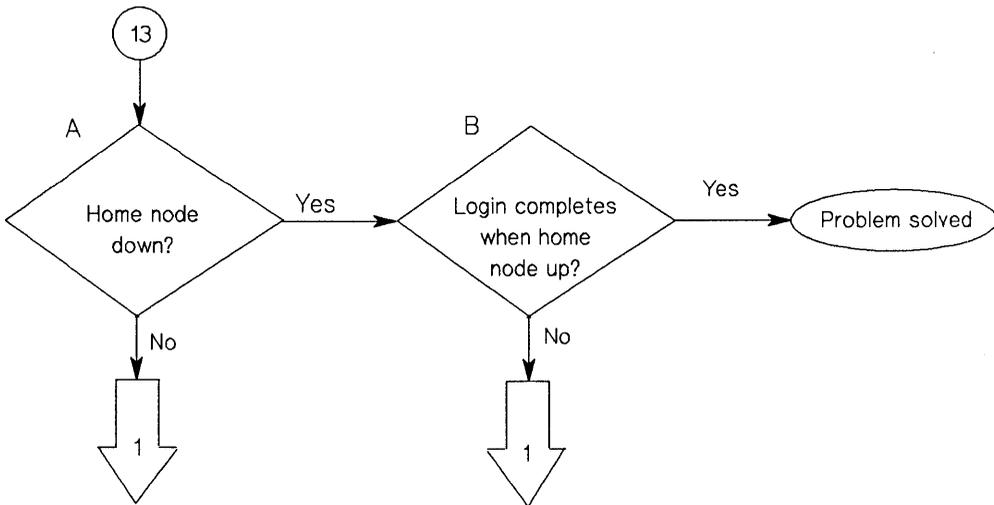
Troubleshoot VHE



Flowchart 12: Initial Steps to Troubleshooting VHE

Initial Steps to Troubleshooting VHE (Flowchart 12)

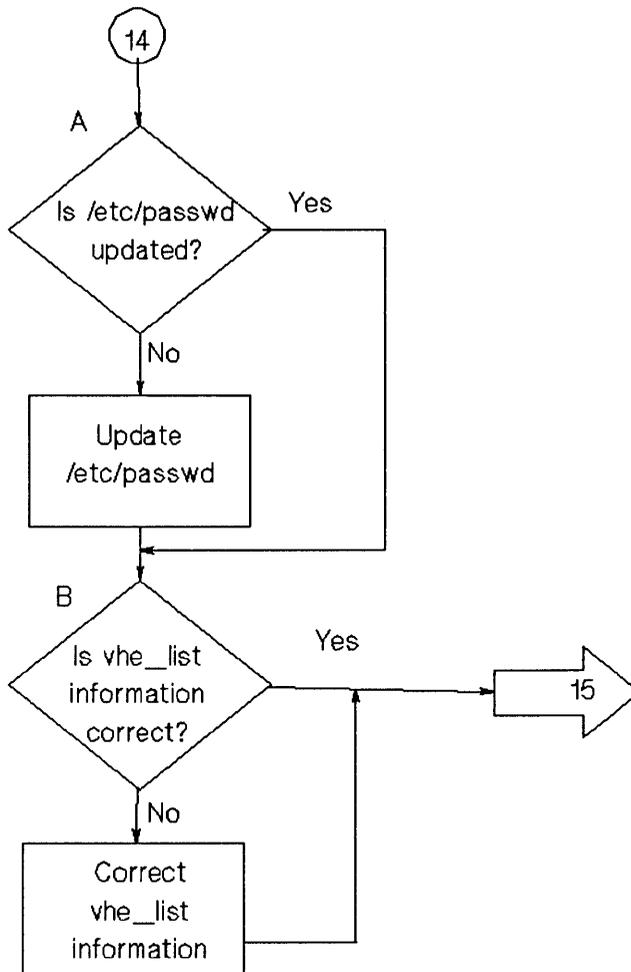
| Question | Yes: Action | No: Action |
|--|---|---|
| A. Are you able to log in? | See B. | See C. |
| B. Are you able to access files on the home node? | No problem. | See D. |
| C. Does the machine hang during login? | See Flowchart 13. | See E. |
| D. Is the home node down? | Retry accessing files 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 YP? | 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 <i>HP-UX System Administrator's Manual/HP 9000 Series 800</i> or <i>HP-UX System Administration Tasks/HP 9000 Series 300</i> manual. |



Flowchart 13: Home Node Goes Down After Mount Done

Home Node Goes Down After Mount Done (Flowchart 13)

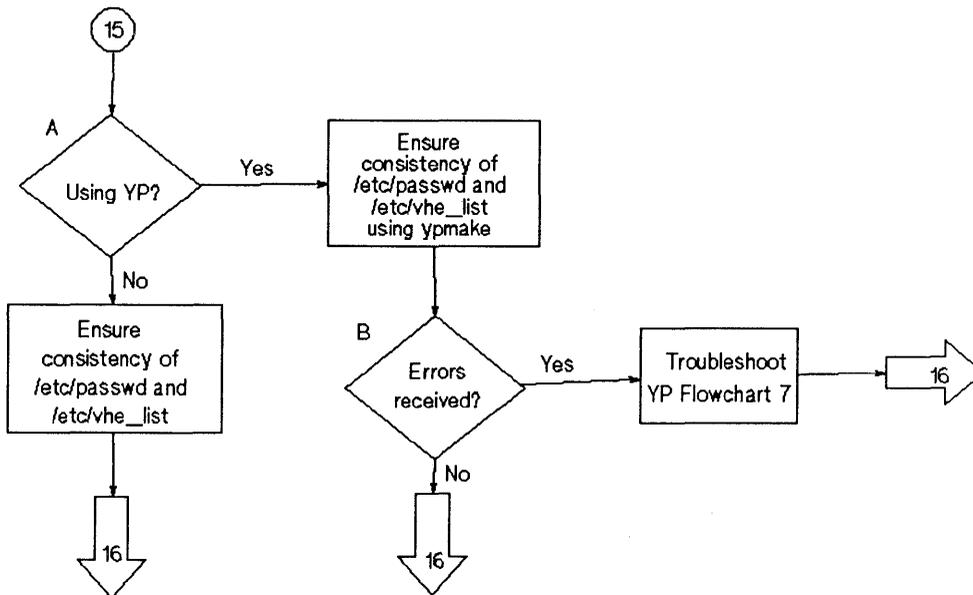
| Question | Yes: Action | No: Action |
|---|---|------------------|
| A. Is the home node down? | Try logging in again once the home node comes up; then see B. | See Flowchart 1. |
| B. Does the login complete once the home node comes up? | Problem solved. | See Flowchart 1. |



Flowchart 14: Checking /etc/passwd and /etc/vhe_list Files

Checking `/etc/passwd` and `/etc/vhe_list` Files (Flowchart 14)

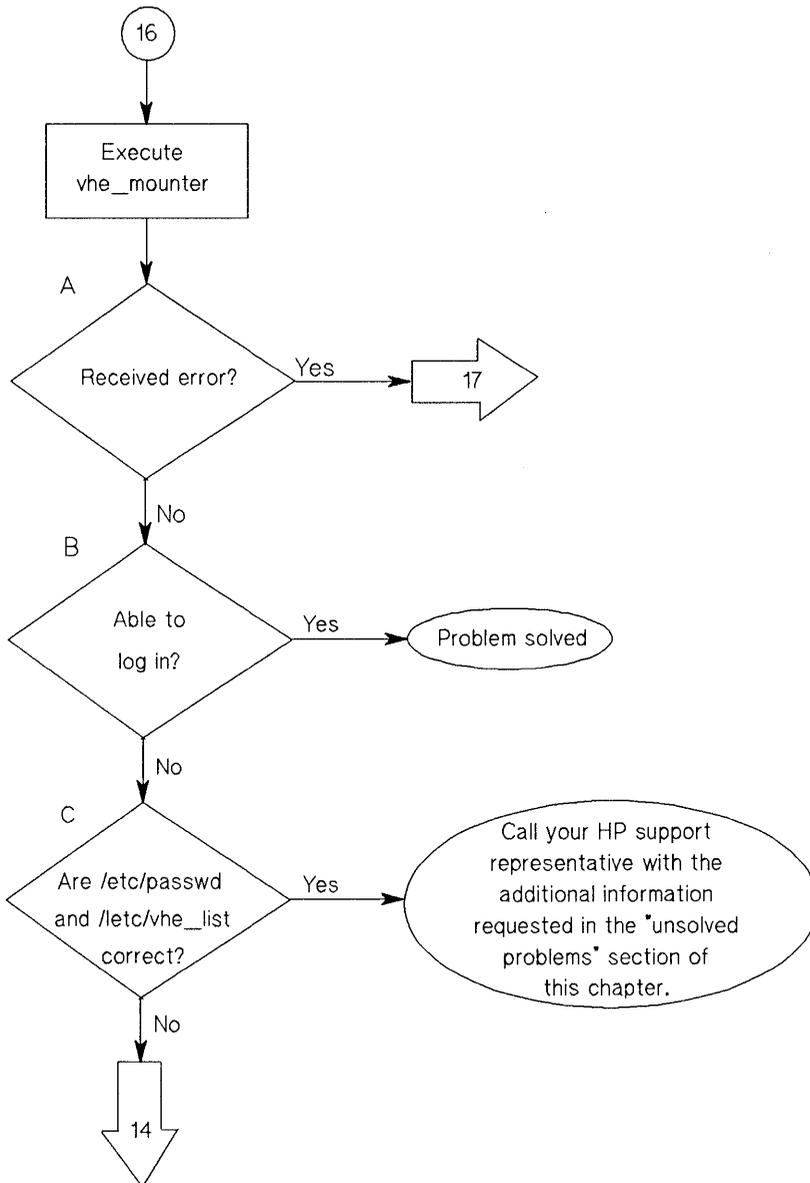
| Question | Yes: Action | No: Action |
|---|-------------------|---|
| A. Is the <code>/etc/passwd</code> file updated to prefix the home directory with the NFS mount point? | See B. | Update the <code>/etc/passwd</code> file as described in the “VHE Configuration and Maintenance” chapter; go to B. |
| B. Is the information in the <code>/etc/vhe_list</code> file correct? | See Flowchart 15. | Correct the <code>/etc/vhe_list</code> 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)

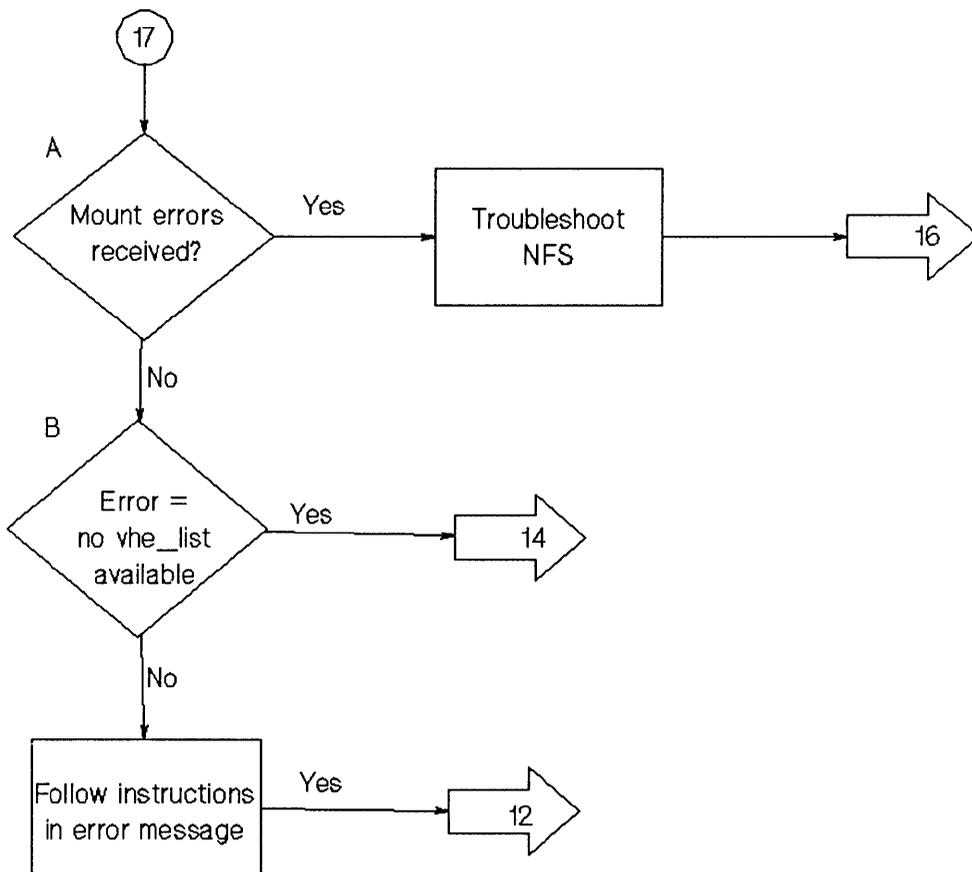
| Question | Yes: Action | No: Action |
|--|--|--|
| <p>A. Are you using Yellow Pages (YP) to ensure consistency of <code>/etc/passwd</code> and <code>/etc/vhe_list</code> information?</p> | <p>Ensure consistency of the <code>/etc/passwd</code> and <code>/etc/vhe_list</code> files on all nodes in the VHE group by executing the following command: <code>/usr/etc/yp/ypmake passwd vhe_list</code> See B.</p> | <p>Ensure consistency of the <code>/etc/passwd</code> and <code>/etc/vhe_list</code> files on all nodes in the VHE group. See Flowchart 16.</p> |
| <p>B. Did you receive any errors when executing <code>ypmake (IM)</code>?</p> | <p>Go to the YP Flowchart 7 and complete troubleshooting steps; then return to VHE Flowchart 16.</p> | <p>Go to Flowchart 16.</p> |



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 <i>vhe_mounter (1M)</i> ? | 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 <i>/etc/passwd</i> and <i>/etc/vhe_list</i> files? | Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter. | See Flowchart 14. |

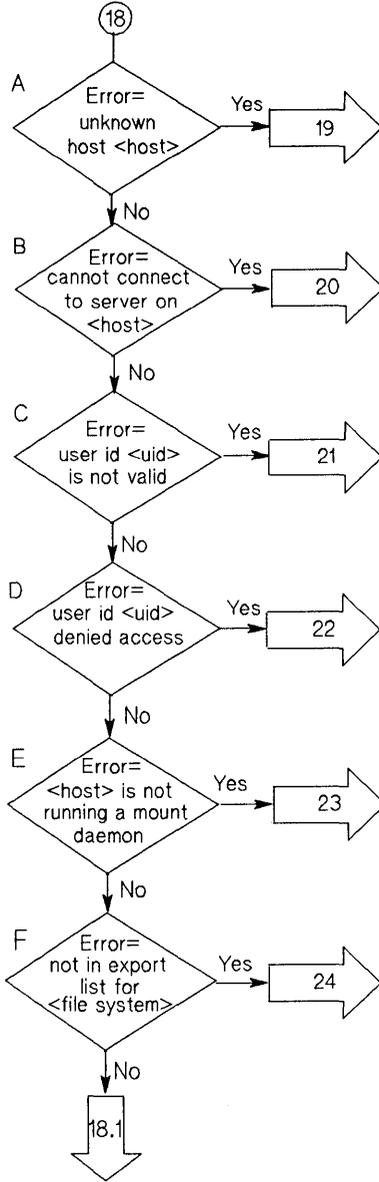


Flowchart 17: Error Message from vhe_mounter

Error Message from vhe_mounter (Flowchart 17)

| Question | Yes: Action | No: Action |
|--|--|---|
| A. Were any mount errors encountered (mount errors begin with <code>mount:</code>)? | Troubleshoot NFS (Flowchart 1); then see Flowchart 16. | See B. |
| B. Does the following error message occur? <code>no vhe_list available</code> | 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. |

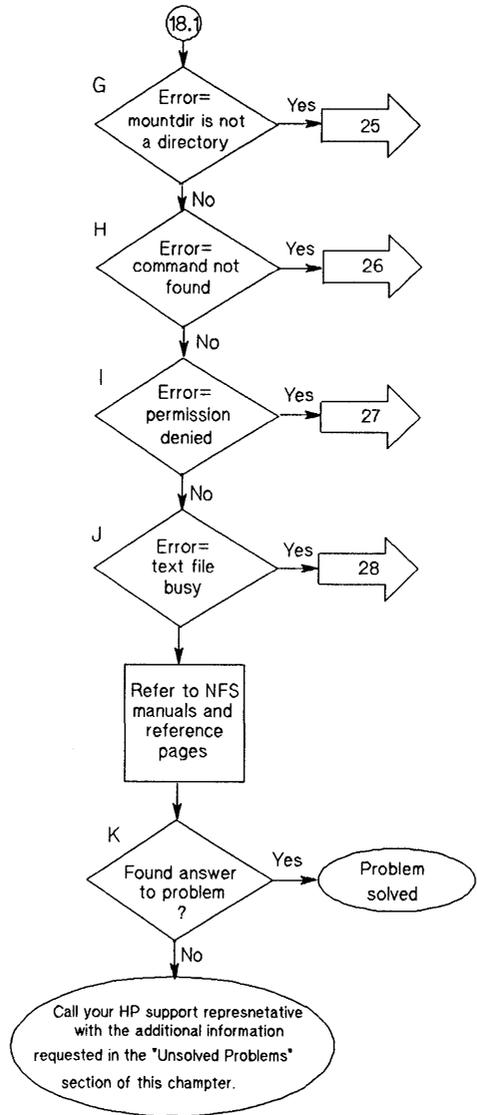
Troubleshoot REX



Flowchart 18: Initial Steps to Troubleshoot REX

Initial Steps to Troubleshoot REX (Flowchart 18)

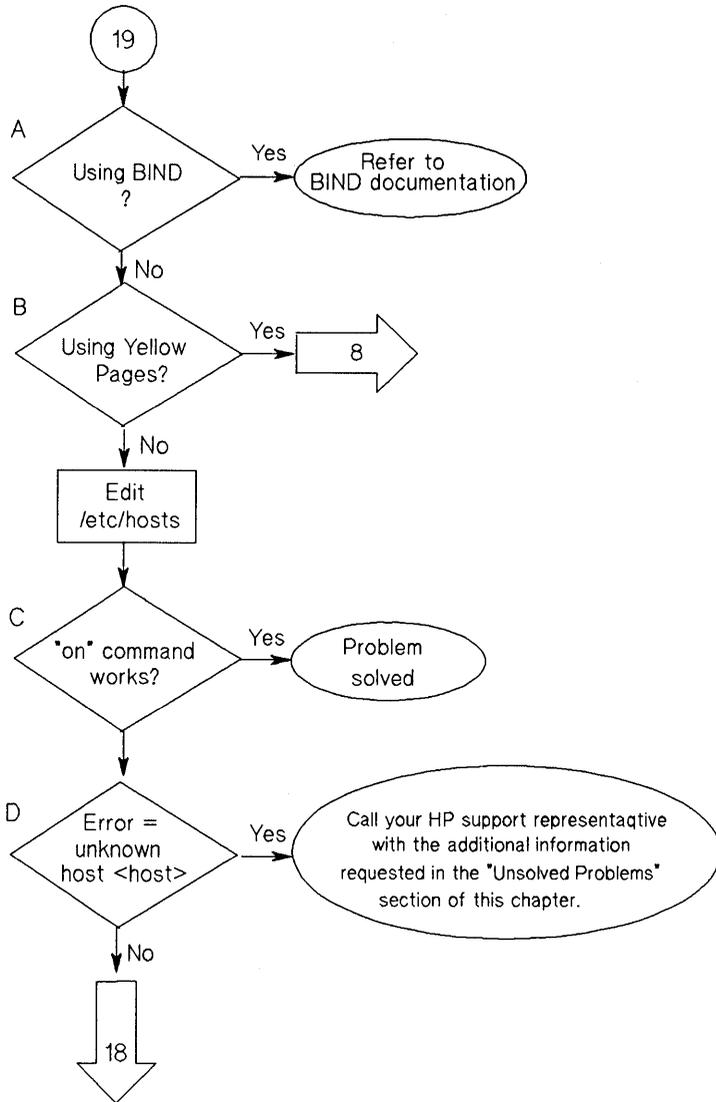
| Question | Yes: Action | No: Action |
|---|-------------------|---------------------|
| <p>A. Does the following error message appear?</p> <p>on: unknown host <host></p> | See Flowchart 19. | See B. |
| <p>B. Does the following error message appear?</p> <p>on: cannot connect to server on <host></p> | See Flowchart 20. | See C. |
| <p>C. Does the following error message appear?</p> <p>on: rexd: user id is not valid</p> | See Flowchart 21. | See D. |
| <p>D: Does the following error message appear?</p> <p>on <server>: rexd: user id <uid>denied access</p> | See Flowchart 22. | See E. |
| <p>E: Does the following error message appear?</p> <p>on: <server>rexd: <host> is not running a mount daemon</p> | See Flowchart 23. | See F. |
| <p>F: Does the following error message appear?</p> <p>on <server>: rexd: not in export list for filesystem</p> | See Flowchart 24. | See Flowchart 18.1. |



Flowchart 18.1: Initial Steps to Troubleshoot REX

Initial Steps to Troubleshoot REX (Flowchart 18.1)

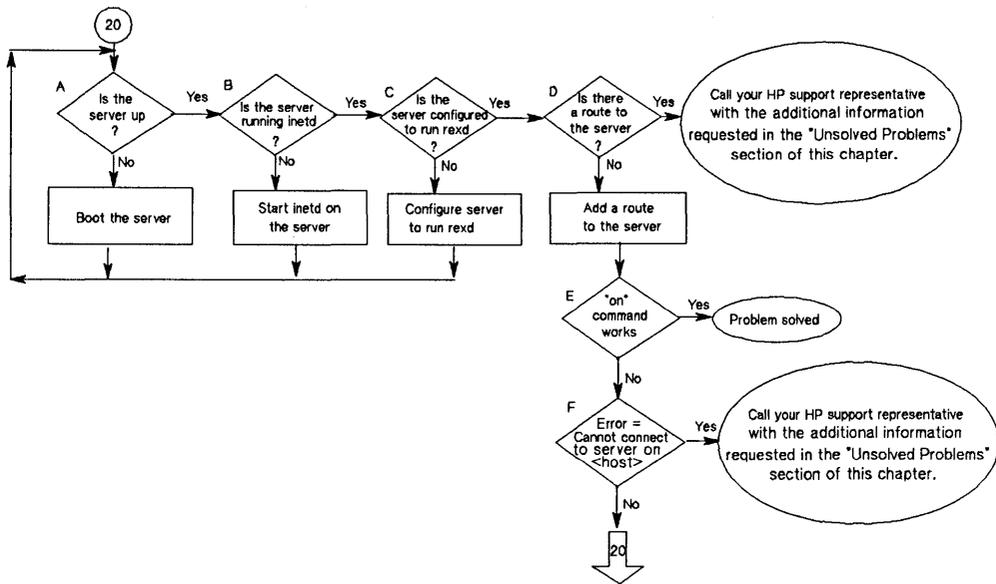
| Question | Yes: Action | No: Action |
|---|-------------------|--|
| <p>G: Does the following error message appear?</p> <p>on <server>: rexd: (<mountdir>) is not a directory</p> | See Flowchart 25. | See H. |
| <p>H: Does the following error message appear?</p> <p>on <server>: rexd: command not found</p> | See Flowchart 26. | See I. |
| <p>I: Does the following error message appear?</p> <p>on <server>: rexd: permission denied</p> | See Flowchart 27. | See J. |
| <p>J: Does the following error message appear?</p> <p>on <server>: rexd: text file busy</p> | See Flowchart 28. | <p>Refer to NFS manuals and reference pages.</p> <p>See K.</p> |
| <p>K: Found answer to your problem?</p> | Problem solved. | <p>Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter.</p> |



Flowchart 19: Unknown Host

Unknown Host (Flowchart 19)

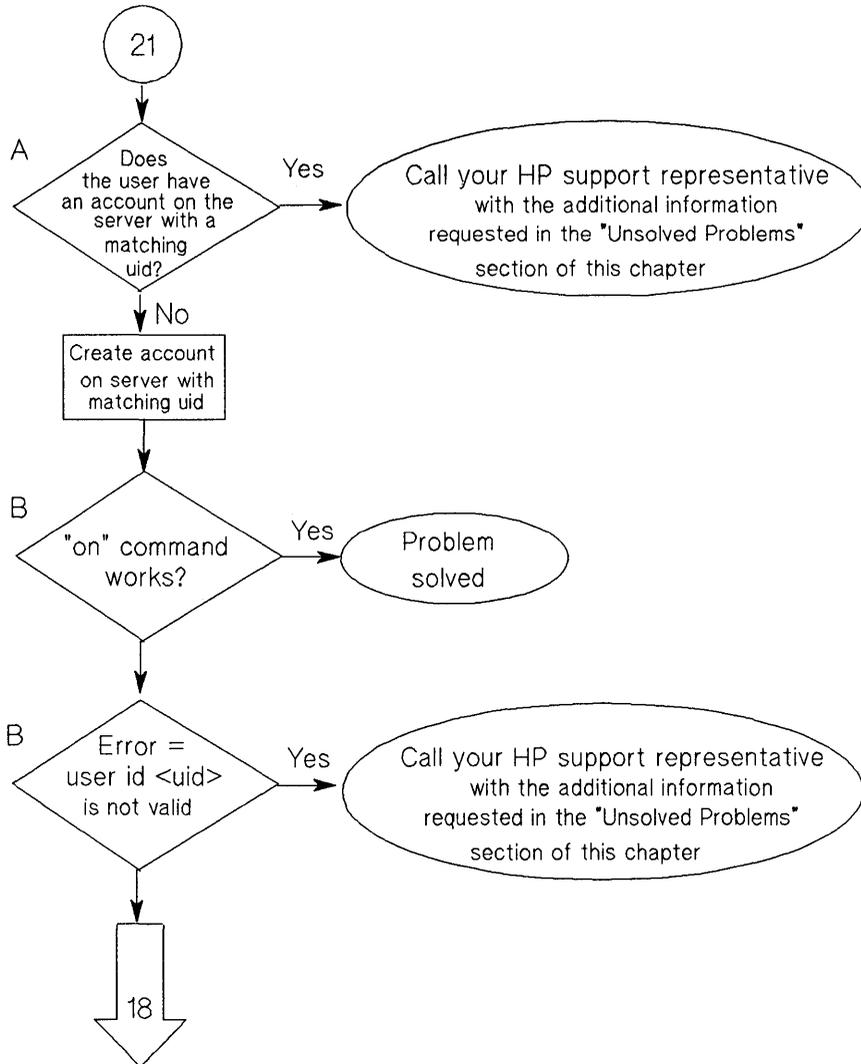
| Question | Yes: Action | No: Action |
|---|---|--|
| A: Is your node using BIND? | Refer to BIND documentation. | See B. |
| B: Is your node using Yellow Pages (YP)? | See Flowchart 8. | Edit <i>/etc/hosts</i> on the client to include the desired remote host. See C. |
| C: <i>on</i> command works now? | 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

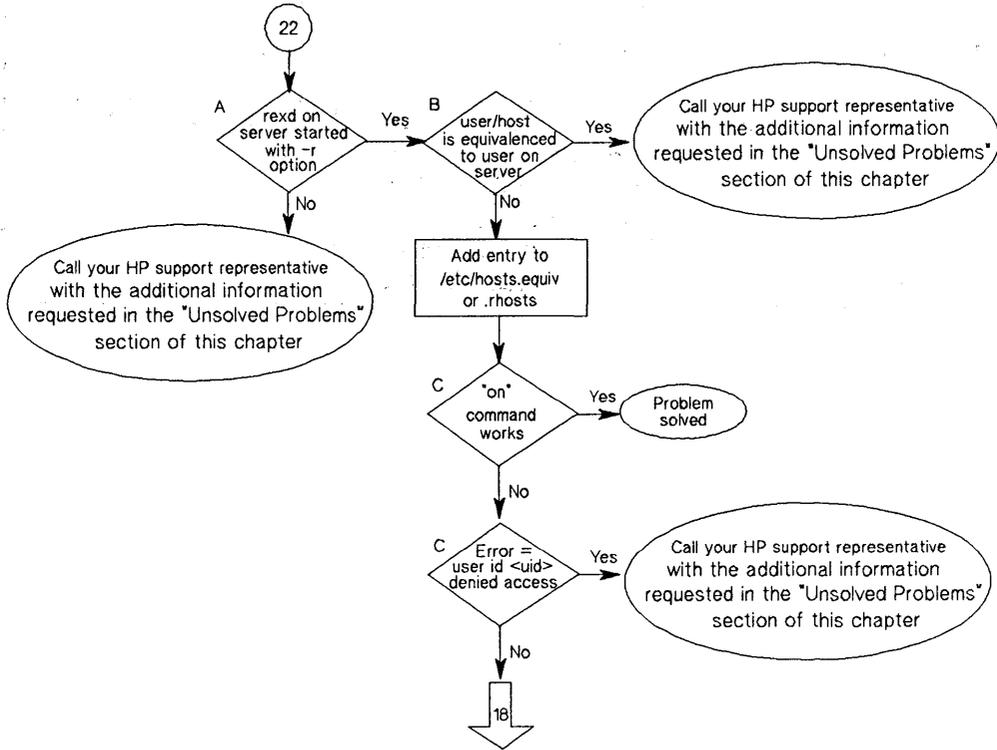
Cannot Connect to REX Server (Flowchart 20)

| 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 <i>inetd</i> on the rex server node. See A. |
| C: Is the rex server configured to run <i>rex</i> ? | See D. | Configure the rex server to run <i>rex</i> by editing <i>/etc/inetd.conf</i> on the rex server, uncommenting the <i>rpc.rexd</i> line, and issuing the <i>inetd -c</i> 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 <i>route(1M)</i> command. See E. |
| E: <i>on</i> 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)

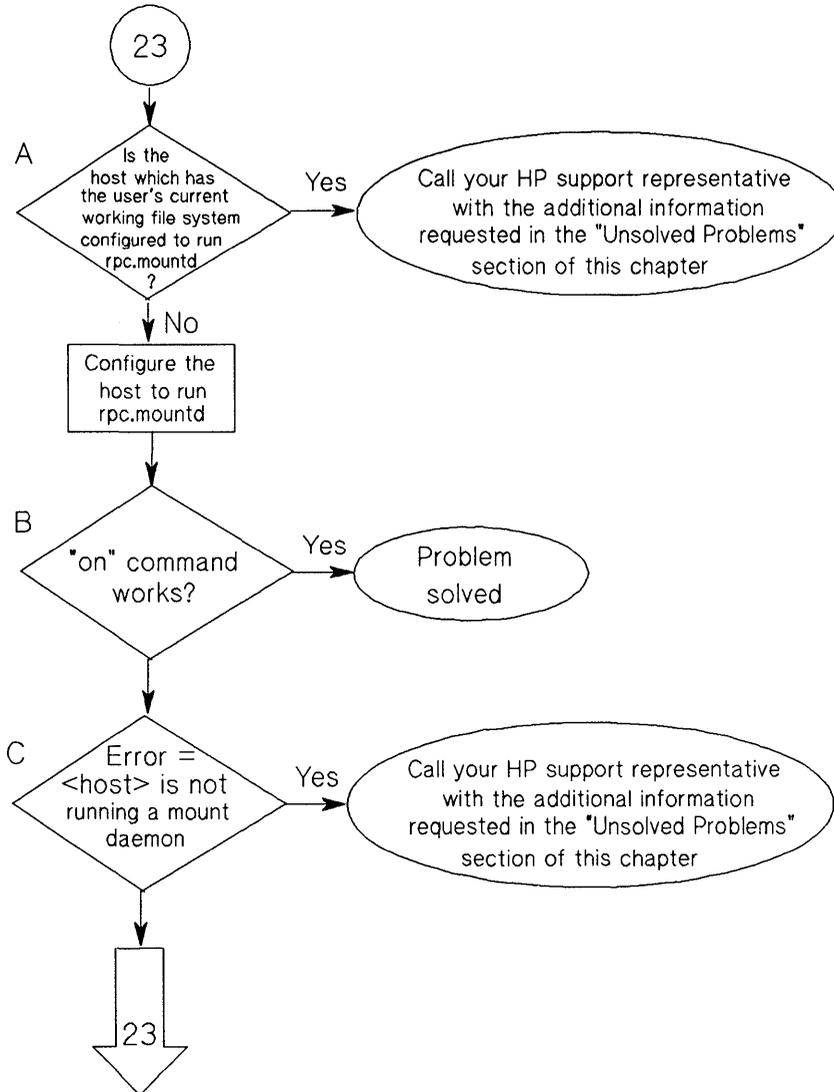
| Question | Yes: Action | No: Action |
|--|---|---|
| A: Does the user have an account on the rex server with a uid which matches the user's uid on the client? | Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter. | Create an account on the rex server for the user with a matching uid. See B . |
| B: <i>on</i> 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)

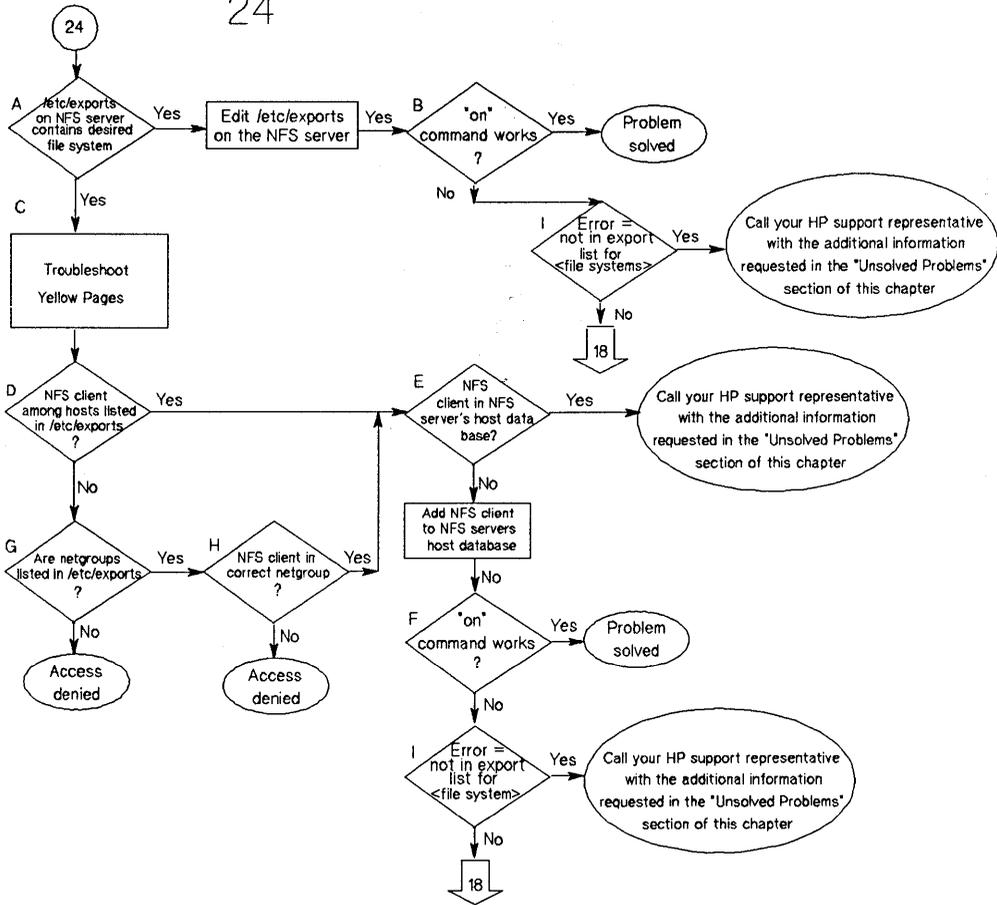
| Question | Yes: Action | No: Action |
|---|---|---|
| A: Rexd or rex server started with <i>-r</i> option? | See B. | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. |
| B: User or client host is equivalenced by entry in <i>.rhosts</i> or <i>/etc/hosts.equiv</i> file? | Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter. | Edit <i>.rhosts</i> or <i>/etc/hosts.equiv</i> file to add an entry for the user or the client host. See C. |
| C: <i>on</i> 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 |
|--|--|--|
| <p>A: Is the host which has the user's current working file system physically mounted configured to run <code>rpc.mountd</code>?</p> | <p>Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.</p> | <p>Configure the NFS server node to run <code>rpc.mountd</code> by editing its <code>/etc/inetd.conf</code> and executing <code>inetd -c</code>.</p> <p>See B.</p> |
| <p>B: <code>on</code> command works now?</p> | <p>Problem solved.</p> | <p>See C.</p> |
| <p>C: Does the following error message appear?</p> <p><code>on: <server>rexd: <host></code> is not running a mount daemon</p> | <p>Call your HP support representative with the additional information requested in the "Unsolved Problems" section of this chapter.</p> | <p>See Flowchart 18.</p> |

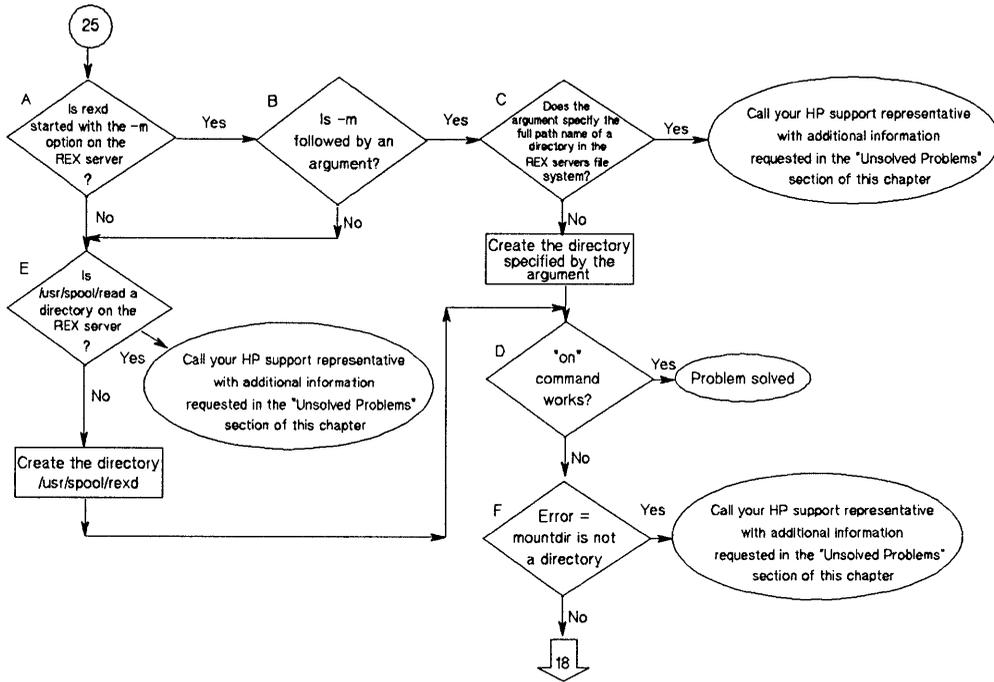


Flowchart 24: REX Server Denied Access through /etc/exports

REX Server Denied Access through /etc/exports (Flowchart 24)

| Question | Yes: Action | No: Action |
|--|---|---|
| A: Does <i>/etc/exports</i> on the NFS server contain desired file system? | Edit <i>/etc/exports</i> on the NFS server to contain the file system rather than the directory. See B. | See C. |
| B: <i>on</i> command works now? | Problem solved. | See I. |
| C: Is your node using Yellow Pages (YP)? | See Flowchart 8. | See D. |
| D: If hosts are listed in the desired <i>/etc/exports</i> 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: <i>on</i> command works now? | Problem solved. | See I. |
| G: Are netgroups found for the desired <i>/etc/exports</i> entry? | See H. | Access for this client is denied. |
| H: Is the client included in a netgroup which is listed in the desired <i>/etc/exports</i> entry? | See E. | Access for this client is denied. |

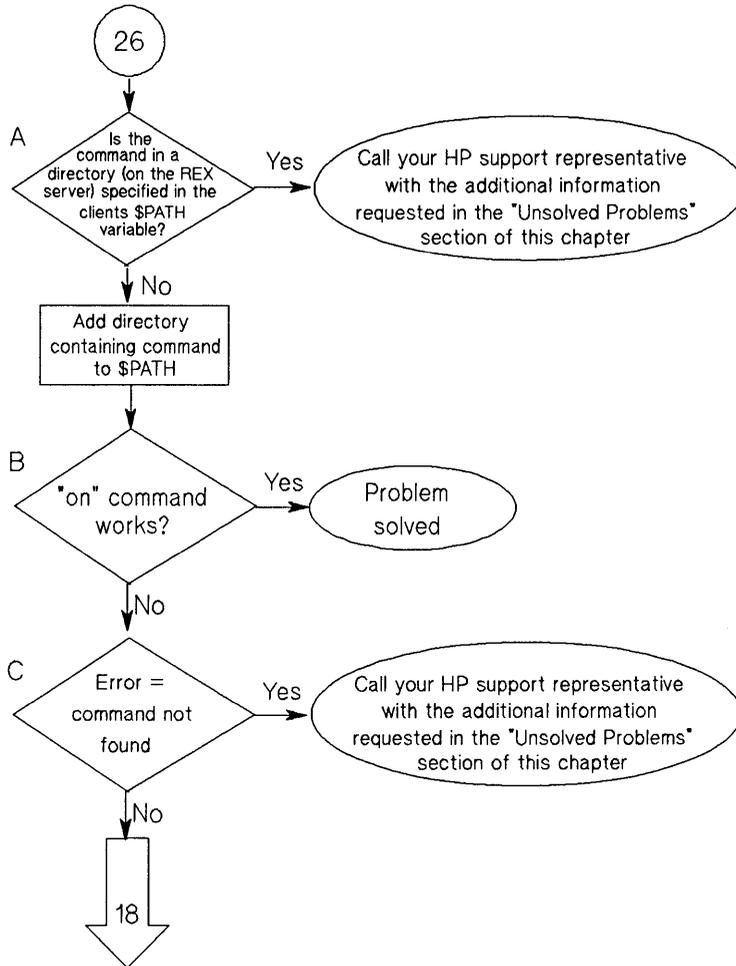
| Question | Yes: Action | No: Action |
|--|---|--|
| <p data-bbox="178 170 450 265">I. Does the following error message appear?</p> <p data-bbox="216 307 471 391">on <server>: rexd: not in export list for file system</p> | <p data-bbox="514 170 807 360">Call your HP support representative with the additional information requested in the “Unsolved Problems” section of this chapter.</p> | <p data-bbox="845 170 1072 194">See Flowchart 18.</p> |



Flowchart 25: Mount Point Not a Directory

Mount Point Not a Directory (Flowchart 25)

| Question | Yes: Action | No: Action |
|--|---|--|
| A: Is <i>rex</i> d on the REX server started with the <i>-m</i> option? | See B. | See E. |
| B: Is <i>-m</i> 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: <i>on</i> command works now? | Problem solved. | See F. |
| E: Is <i>/usr/spool/rexd</i> 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 <i>/usr/spool/rexd</i> . See D. |
| F: Does the following error message appear? on <server>:rex: (<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. |

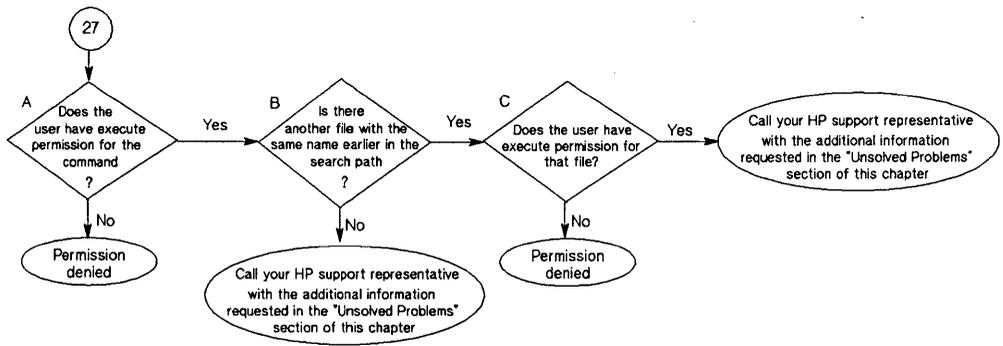


Flowchart 26: Command Not Found

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 <i>\$PATH</i> 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 <i>\$PATH</i> variable. See B. |
| B: <i>on</i> 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. |

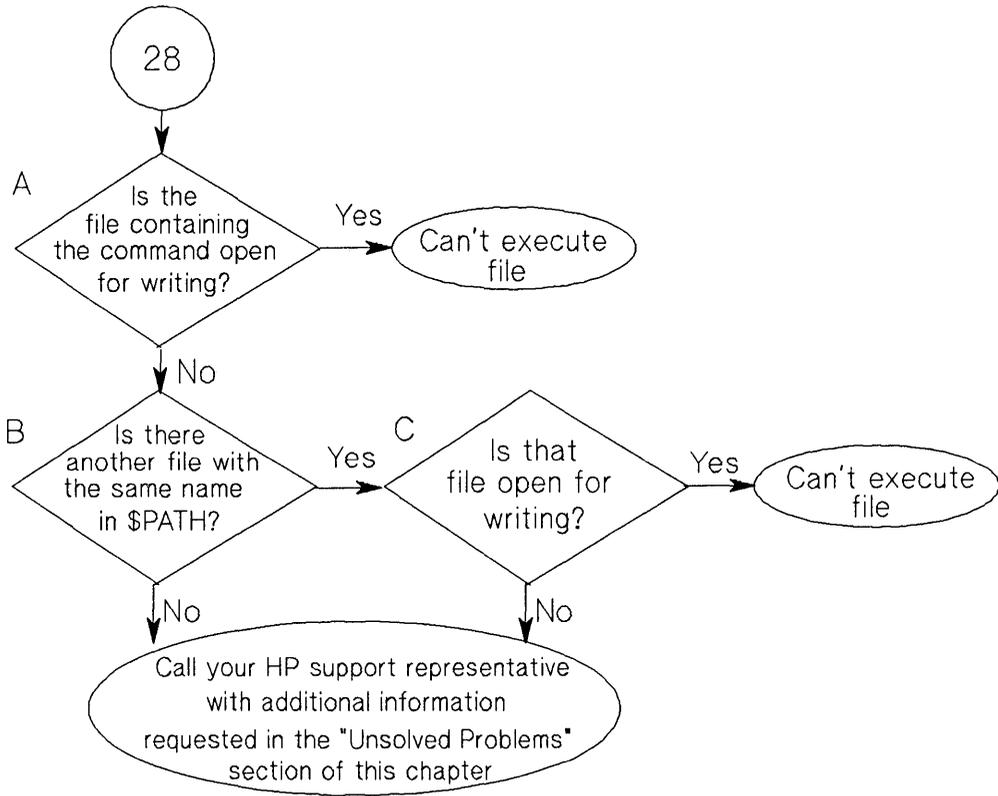
27



Flowchart 27: Permission Denied

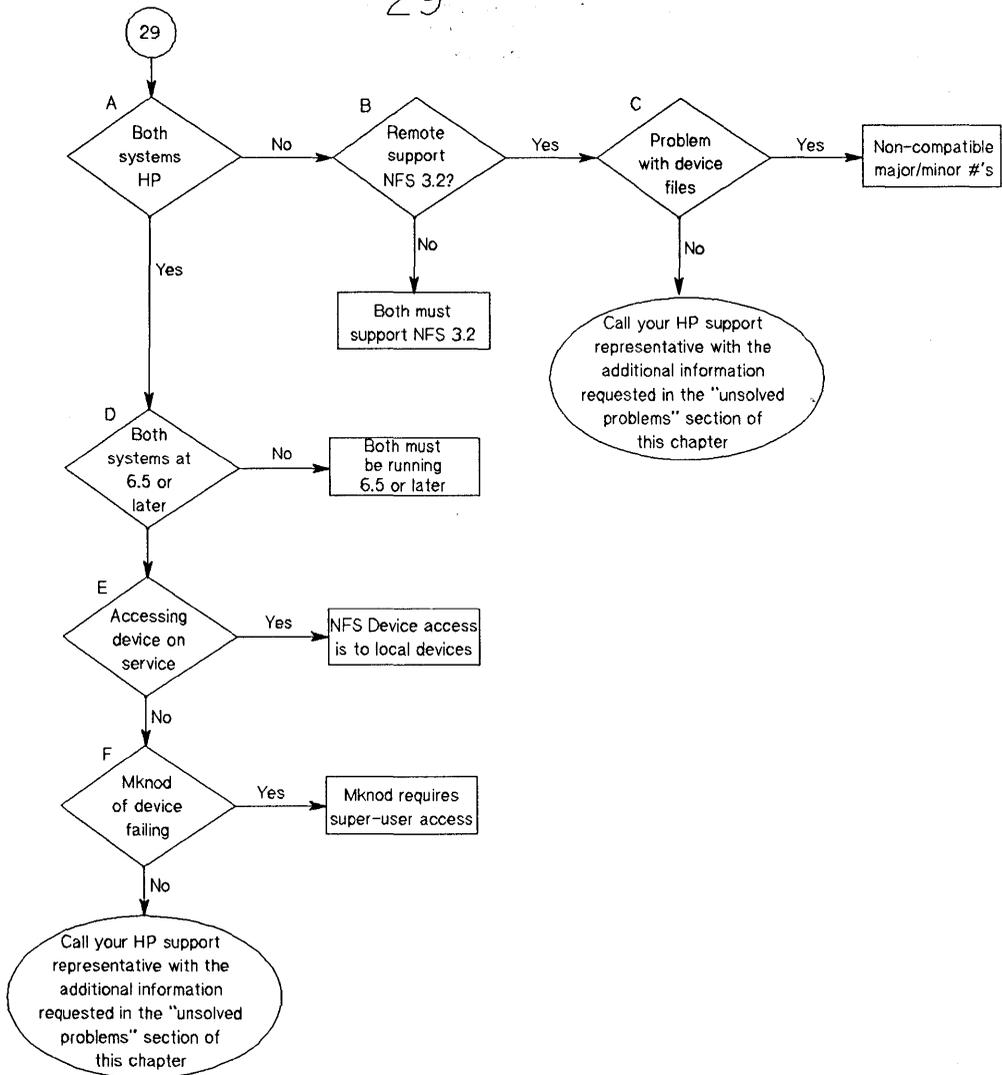
Permission Denied (Flowchart 27)

| 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 <i>\$PATH</i> 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 <i>\$PATH</i> 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 “-noders” 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 or later? | See E. | Both systems must be running HP-UX release 6.5 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 <code>mknod(1M)</code> a device file failing? | Creation of device files requires super-user 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 Networking Operation | Local HP-UX Operation |
|---|--|
| setaclentry(3) library routine | |
| Is not supported over NFS. | Is supported locally. |
| Super-user Permission | |
| <p>The super-user UID 0 is mapped to -2 by default.</p> <p>Anything requiring super-user permission may not work over NFS. For example, a super-user may not be able to perform the following tasks.</p> <ul style="list-style-type: none"> ■ Link and unlink directories ■ Alter directories such as /, /etc, and /bin ■ Use chmod to set sticky or setuid bits ■ <i>mknod</i> of device files | <p>Super-user has permission to perform any operation locally (by definition).</p> |
| System Time | |
| <p>Commands that access clocks on different systems may not provide consistent times since system clocks differ.</p> <p>For example, if you give the <i>utime(2)</i> command a <i>NULL</i> pointer for the times value, the following process occurs:</p> <ol style="list-style-type: none"> 1. The system sets the access time and modification time according to the client node clock. | <p>Commands that access clocks on the local system provide consistent times.</p> |

| HP NFS Services Networking Operation | Local HP-UX Operation |
|--|--|
| System Time continued... | |
| <p>2. It then sends these times over to the server which changes the inode to reflect the new access and modification times.</p> <p>3. The server node identifies the change in the inode and thus, modifies the inode's status change time according to its own clock.</p> <p>The result is a high probability of differing times between the server's access and modification times versus its status change time.</p> <p>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.</p> | |
| Unlinking | |
| <p>The server does not keep state information and does not know if a process has a file open.</p> <p>■ 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.</p> | <p>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.</p> |

| HP NFS Services Networking Operation | Local HP-UX Operation |
|---|---|
| Unlinking <i>continued...</i> | |
| <ul style="list-style-type: none"> ■ 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. | |
| yppasswd(1) Command | passwd(1) Command |
| <p>This command does not have a password aging feature.</p> <p>The super-user must know the current password to change another user's password.</p> <ul style="list-style-type: none"> ■ A password must contain at least five characters if it includes a combination of either: uppercase and lowercase letters, numeric or special characters. ■ A password must contain at least four characters if it includes a combination of uppercase letters, lowercase letters, and numeric characters. ■ A password must contain at least six characters if it includes only monospace letters. | <p>This command has a password aging feature.</p> <p>Super-user does not have to know the password to change another user's password.</p> <ul style="list-style-type: none"> ■ Each password must have six or more characters: at least two alpha characters and at least one numeric or special character. ■ 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. |

| HP NFS Services Networking Operation | Local HP-UX Operation |
|--|---|
| pathconf/fpathconf | |
| <p>The following variables for the <i>pathconf/fpathconf</i> system calls are not supported over NFS:</p> <p><u>PC_CHOWN_RESTRICTED</u> variable</p> <p><u>PC_LINK_MAX</u> variable</p> <p><u>PC_NAME_MAX</u> variable</p> <p><u>PC_NO_TRUNC</u> variable</p> <p><u>PC_PATH_MAX</u> variable</p> <p>The following variables for the <i>pathconf/fpathconf</i> system calls return local information over NFS.</p> <p><u>PC_MAX_CANON</u> variable</p> <p><u>PC_MAX_INPUT</u> variable</p> <p><u>PC_VDISABLE</u> variable</p> <p>The following variable for the <i>pathconf/fpathconf</i> systems calls is supported over NFS:</p> <p><u>PC_PIPE_BUF</u> variable</p> | <p>All variables are supported locally for the <i>pathconf/fpathconf</i> system calls:</p> <p><u>PC_CHOWN_RESTRICTED</u> variable</p> <p><u>PC_LINK_MAX</u> variable</p> <p><u>PC_NAME_MAX</u> variable</p> <p><u>PC_NO_TRUNC</u> variable</p> <p><u>PC_PATH_MAX</u> variable</p> <p><u>PC_MAX_CANON</u> variable</p> <p><u>PC_MAX_INPUT</u> variable</p> <p><u>PC_VDISABLE</u> variable</p> <p><u>PC_PIPE_BUF</u> variable</p> |

Migrating from RFA to NFS

When using networks consisting of all HP systems, the Remote File Access (RFA) service provides distributed file access among Series 300 and 800 computers.

Use this appendix if you wish to translate your RFA applications to NFS applications.

Why Migrate to NFS Services?

Using NFS Services 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 YP provides centrally administered databases.

Similarities

HP NFS and RFA have the following similarities.

- No remote device access
- Not all UNIX¹ semantics are fully supported

Differences

Refer to the following table for a list of differences between HP NFS and RFA.

| NFS Services | RFA |
|--|--|
| You can run <i>setuid</i> programs accessing data on remote file systems. | You cannot run <i>setuid</i> 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 super-user can perform remote NFS mounts. | All users can establish access to remote file systems. |
| You can centrally administer your databases using YP service. | You have no centrally administered database. |
| All users with read access to the mount point can read the remote file system. | Only users performing <i>netunam</i> 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. |

(1) UNIX (R) is a U.S. registered trademark of AT&T in the U.S.A. and other countries.

| NFS Services | RFA |
|---|---|
| One mount gives you access to only one file system. | One <i>netunam</i> 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 only minor modifications.

- You must perform a remote mount of a file system or directory either
 - as part of the script or
 - before executing the script.

Since super-user must execute mounts, the script must be *setuid root* if the mount is performed as part of the script.

Note Having *setuid root* scripts is a potential security problem.

If the script's owner does not have super-user permissions, the super-user 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.

- If RFA is **not** being used, remove any 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 above in “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*, and you want to 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*.

1. Remove the network special file */net/systemB*.
2. Create the directories */net/systemB* and */net/systemB/project*.

mount systemB:/project /net/systemB/project

The path name, therefore, remains the same.

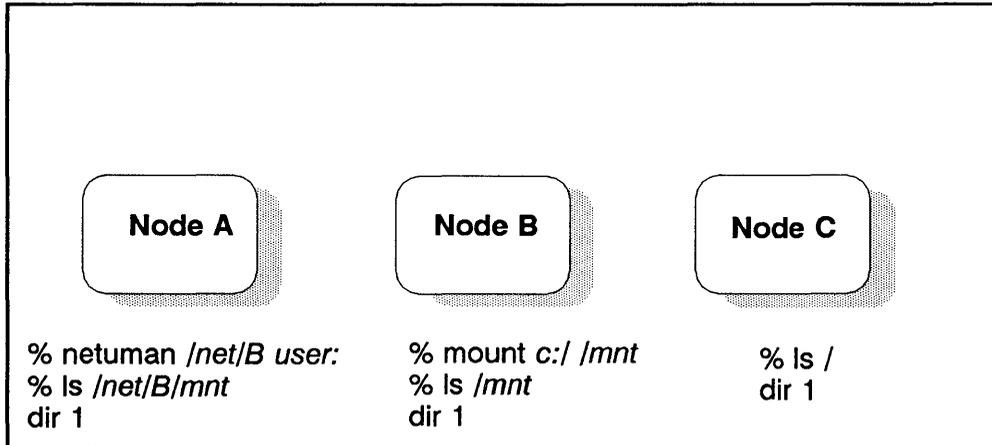
Note For RFA, access to the remote system is via a network special file. Creating an NFS mount point with the same name as the network special file for the remote system could cause confusion. Problems will not occur if

- the system does not use RFA and 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.

RFA through NFS

RFA functions are operational through NFS. For example, an RFA node can access a remote directory by going through an NFS client. Note, however, **NFS functions cannot operate through RFA** because NFS parses the path names differently than RFA. Therefore, an NFS node cannot access a remote node by going through RFA.



EXAMPLE: RFA remote mount through an NFS client

1. Node B performs an NFS mount to Node C.
2. Node A then performs a netunam to Node B.
3. Node A lists (ls) the contents of Node C's root directory.

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 |
|----------------------------|--|
| CDF | Context Dependent File: 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. |
| Cnode | Any node operating in an HP-UX cluster environment, including diskless nodes and the root server. |
| Diskless Cnode | A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a file system directly attached. |
| Homogeneous Cluster | A diskless cluster composed of nodes of only one computer architecture (e.g., Series 300 only). |
| Mixed Cluster | Diskless cluster consisting of a Series 800 diskless server and Series 300/800 diskless clients. |
| Root Server | The only node in an HP-UX cluster that has file systems directly attached to it. |

NFS Configuration and Maintenance

Configure

If you configure NFS on the root server, you must also configure NFS on all clients in the cluster. If the root server does not have NFS configured, then none of the clients can use NFS.

Daemons

- The *nfsd(1M)* daemon should be running on the root server if it is servicing NFS requests. Any *nfsd(1M)* daemons running on client cnodes are ignored.
- The *mountd(1M)* daemon should be running on the root server if servicing NFS requests. Any *mountd(1M)* daemons running on client cnodes are ignored.
- The *biod(1M)* daemon must 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 root server host's name as the node name specified in the *mount(1M)* 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*.

**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 root server's clock. Therefore, clock skew problems exists only if the root server's clock is different from other NFS servers.

YP Configuration and Maintenance

HP recommends that you execute *ypserv(1M)* only on the root server

- for better performance and
- to ensure the root server is the only YP server for that cluster.

Troubleshooting

You can troubleshoot NFS specific problems from the root server.

- If trying to mount an NFS file system, ensure you are using the root server's host name as the node specified in the *mount(1M)* 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 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 Yellow Pages. If you wish to review the normal use of passwords with Yellow Pages, see the “YP Configuration and Maintenance” section in this manual. If you require additional information on the secure password file, see *passwd(4)* 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 YP database will not contain encrypted passwords if you use this */etc/passwd* file to build your Yellow Pages (YP) password database. This prevents non-privileged users from reading your passwords, as anyone with access to YP commands such as *yycat* or YP library routines such as *yp_first* and *yp_next* can read the YP 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 Yellow Pages (YP) database.

Note A password in the */.secure/etc/passwd* file takes precedence over the password stored in Yellow Pages (YP).

If you wish to hide the encrypted passwords in your HP systems and wish to continue to use the YP password database to maintain other information kept on the password file, you can do the following:

- Build your YP password database on the HP YP master server using a password file that does not contain encrypted passwords (e.g. uses "*" in the password field).
- On an HP YP client, maintain a copy of the secure password file so the passwords in that file will be used at login.

and/or

- On an HP or non-HP Yellow Pages client, maintain the encrypted password in the */etc/passwd* file through a YP escape.

EXAMPLE:

```
+username:encrypted_passwd:::::
```

Glossary

A

- 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.

B

- 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.

C

- 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, YP client, or YP server. (A YP server **must** also be configured as a YP 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.
- Cnode** Any node operating in an HP-UX cluster environment, including diskless nodes and the root server.

D

- Daemon** Background programs that are always running, waiting for a request to perform a task.
- Diskless Cnode** A node in an HP-UX cluster that uses networking capabilities to share file systems, but does not have a file system directly attached.

E

Escape Sequence (YP)

Characters used within files to force inclusion and exclusion of data from YP 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.

F

File System

A directory structure used to organize files.

G

GID

A value that identifies a group in HP-UX.

Global (YP)

A means of access in which the system always reads YP maps rather than the local ASCII files.

H

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.

I

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.

Interrupt-able 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.)

K

Key (YP) A string of characters (no imbedded blanks or tabs) that indexes the values within a YP 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.

L

Local (YP) A means of access in which the system first reads the local ASCII file. If it encounters an escape sequence, it then accesses the YP databases.

M

- Map (YP)** A file consisting of logical records; a search key and related value form each record. YP clients can request the value associated with any key within a map.
- YP map** is synonymous with **YP database**.
- Map Nickname (YP)** A synonym for the YP map name when using certain YP commands.
- Master Server (YP)** The node on which one or more YP maps are constructed from ASCII files. These maps are then copied to the YP slave servers for the YP 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.

N

- Netgroup** A network-wide group of nodes and users defined in */etc/netgroup*.
- NFS** Network File System.
- 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.
- Node** A computer system that is attached to or is part of a computer network.

P

- Propagate** To copy maps (data) from one YP server to another.
- Protocol** The rules and steps by which servers and clients exchange data and control information.

R

- 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.
- Root Server** The only node in an HP-UX cluster that has file systems directly attached to it.

S

- 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, YP client, or YP server. (A YP server **must** also be configured as a YP client.)

| | |
|--------------------------|--|
| Slave Server (YP) | A node that copies YP maps from the YP master server and then provides YP 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 | 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 YP, the information contained in a YP map is consistent among all YP servers within a given YP domain (i.e., is not in the process of being updated). |

U

| | |
|----------------|---|
| 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 <i>mount(IM)</i> command. |
| Update | The HP-UX command that installs software onto the system. |

V

Value (YP)

A unit of information stored in YP 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.

X

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.

Y

Yellow Pages (YP)

An optional network service composed of databases (maps) and processes that provide YP clients access to the maps. The YP service enables you to administer these databases from one node.

YP may or may not be active; check with your system administrator.

YP Client

■ A node that requests data or services from YP servers.

■ A YP process that requests other YP processes to perform operations.

Note: A YP client can also be configured as any combination of a YP server, NFS client, or NFS server. (A YP server **must** also be configured as a YP client.)

YP Database

See “Map (YP).”

YP Domain

A logical grouping of YP maps (databases) stored in one location. YP domains are specific to the YP network service and are not associated with other network domains.

YP Map

See “Map (YP).”

**YP
Password**

The password for a user's login ID that exists in the YP *passwd* map. The YP password is the same one as the user password, but is administered through the YP.

You do not have to have a YP password to access the YP databases.

YP Server

- A node that provides data (maps) or services to other nodes (YP clients) on the network using YP.
- A YP process that performs operations as requested by other YP processes.

Note: A YP server **must** also be configured as a YP client. It can also be configured as an NFS server, NFS client, or both.

Index

A

Adding Computers, 3-10
Automatic Mounts, 4-40

B

Binding, 2-2

C

Clients

NFS, 2-1

Unmount File Systems,
4-45

YP, 2-14

Clock Skews, 4-50

Clusters, C-1

CDFs, C-1, C-3

Clock Skews, C-3

Daemons, C-2

Mounts, C-2

NFS Configuration, C-2

NFS Maintenance, C-2

Troubleshooting, C-3

Unmounts, C-2

YP Configuration, C-3

YP Maintenance, C-3

Commands

domainname(1), 7-12

makedbm(1M), 7-12

on(), 5-1, 5-3

YP, 7-12

ypbind(1M), 7-12

ypcat(1), 7-12

ypinit(1M), 7-12

ypmatch(1), 7-13

yppasswd(1), 7-13

yppasswdd(1M), 7-13

yppoll(1M), 7-13

ypserv(1M), 7-13

ypwhich(1), 7-14

ypxfr(1M), 7-14

Configuration, 4-15

NFS, 4-1, 4-11

YP, 7-14

Configuration Files

/etc/checklist, 4-40

/etc/exports, 4-6, 4-25

/etc/hosts, 4-22

/etc/inetd.conf, 4-8, 4-18

/etc/netgroup, 4-23

/etc/netnfsrsc, 4-6, 4-16, 4-31

/etc/rpc, 4-6

/usr/adm/inetd.sec, 4-8, 4-20

NFS, 4-6

Configure Kernel, 3-9

crontab(1), 7-29

D

Daemons

Clusters, C-2

inetd(1M), 4-9, 4-18

NFS, 4-8

pcnfsd(1M), 4-9

- portmap(1M)*, 4-9
- rexid(1M)*, 5-1
- Device Files, 2-6
- Diagnostics, 5-10, 5-12
- Documentation
 - Contents, 1-2
 - Conventions, 1-5
 - Overview, 1-1—1-2
- domainname(1)*, 7-12

E

Editing

- /etc/checklist*, 4-40
- /etc/exports*, 4-25
- /etc/hosts*, 4-22
- /etc/inetd.conf*, 4-18
- /etc/netfsrc*, 4-31
- /etc/netgroup*, 4-23
- /etc/netgroups*, 7-8
- /etc/netnfsrc*, 4-16
- /usr/adm/inetd.sec*, 4-20
- Environment Simulation, 5-5—5-6
- EREMOTE, 9-14
- Errnos
 - EREMOTE, 9-14
 - ESTALE, 9-14
- Error Messages, 9-13
 - on* command, 5-5, 5-10
 - rexid*, 5-11
- Escape Sequences, 7-7
- ESTALE, 9-14
- Export File Systems, 4-25
- External Data Representation, 2-10

F

- fcntl()*, 6-1
- fg* Mount Defaults, 4-37
- File Access, 2-4

- Limiting, 4-22—4-23, 4-25
- Removing, 4-45

File Systems

- Automatic Mounts, 4-40
- Availability of, 4-22—4-23
- Manual Mounts, 4-42
- Mounting of, 4-34
- Preventing Access, 4-47
- Removing Access, 4-45
- Unmount, 4-45

Files

- NFS Configuration, 4-6

- Flowchart Formats, 9-15

G

- GIDs, Setting of, 4-14
- Global Maps, 7-5

H

- Hard Mounts, 4-34, 4-39—4-40, 4-42
- \$HOME/.rhosts*, 5-8, 7-11
- HP-UX Clusters, C-1

I

- inetd(1M)*, 4-9, 4-18
- Installation, 3-1
 - Adding Computers, 3-10
 - Configure Kernel, 3-9
 - Preparing the System, 3-5
 - Recompile Programs, 3-11
 - Software, 3-6
 - update, 3-7
- int* Mount Defaults, 4-39

L

- Local Maps, 7-5
- lockf()*, 6-1
- Locking Protocol, 6-3—6-4
- Log Files, 7-43

M

- Maintenance
 - NFS, 4-45
 - YP, 7-33
- makedbm(1M)*, 7-12
- Manual Mounts, 4-42
- Maps, 2-14
 - Global, 7-5
 - Local, 7-5
 - Maintenance, 7-34
 - Manual Modifications, 7-35
 - Modification, 7-34
 - Non-standard, 7-35, 7-44
 - Propagation, 7-29
- Master Server
 - Changing of, 7-40
- Master Servers
 - Automatic Start, 7-18
 - Configuration, 7-15
 - Manual Start, 7-18
 - Security, 7-16
- MAX_INPUT*, A-6
- Memory, 4-5
- Mount Defaults
 - fg*, 4-37
 - hard*, 4-39
 - int*, 4-39
 - port*, 4-39
 - setuid*, 4-38
 - wsize*, 4-39
- Mount Options
 - noauto*, 4-37

- nosuid*, 4-37
- ro*, 4-39
- soft*, 4-38
- mount(1M)*, 4-42
- Mounts
 - /etc/checklist*, 4-40
 - Automatic, 4-40
 - Guidelines, 4-35
 - Hard, 4-34, 4-39—4-40, 4-42
 - Manual, 4-42
 - Options, 4-37
 - Soft, 4-35, 4-38, 4-40, 4-42
 - Wide Area Networks, 4-39

N

- Named Pipes, 2-5
- Netgroups, 4-23, 7-8
- Network Lock Manager, 2-11, 6-1, 6-3, 6-5
- Network Memory, 4-5
- Network Status Monitor, 6-5—6-6
- NFS
 - Daemons, 4-8
 - Servers, 4-10
- NFS Client
 - Configuration, 4-31
- NFS Clients, 2-1
- NFS Configuration, 4-1, 4-11
 - Clients, 4-31
 - Clusters, C-2
 - Edit */etc/checklist*, 4-40
 - Edit */etc/exports*, 4-25
 - Edit */etc/hosts*, 4-22
 - Edit */etc/inetd.conf*, 4-18
 - Edit */etc/netfsrc*, 4-31
 - Edit */etc/netgroup*, 4-23
 - Edit */etc/netnfsrc*, 4-16
 - Edit */usr/adm/inetd.sec*, 4-20
 - Files, 4-6
 - Guidelines, 4-5
 - Memory, 4-5

- Reboot System, 4-30, 4-44
- Security, 4-18, 4-20
- Servers, 4-16
- Set UIDs and GIDs, 4-14
- NFS Maintenance, 4-45
 - Clock Skews, 4-50
 - Clusters, C-2
 - Remove NFS File Access, 4-45
 - Update Software, 4-49
- NFS Servers, 2-1
 - Configuration, 4-16
 - Security, 4-18, 4-20
- NFS Services
 - Overview, 2-1, 2-3
- NFS to RFA
 - Changing Scripts, B-4
- NFS Troubleshooting, 9-11, 9-17
- noauto* Mount Options, 4-37
- Non-standard Maps, 7-35
- nosuid* Mount Options, 4-37

P

- PC_CHOWN_RESTRICTED*, A-6
- PC NFS Servers, 4-16, 4-31
- pcnfsd(1M)*, 4-9
- port* Mount Defaults, 4-39
- portmap(1M)*, 4-9
- Program Recompiling, 3-11
- Propagate Maps, 7-29

R

- Reboot System, 4-30, 4-44
- Recompile Programs, 3-11

- Remote Execution Facility, 2-8, 5-1, 5-3, 5-5, 5-7, 5-9, 5-11
- Remote File Access, 2-4
 - Limiting, 4-22—4-23, 4-25
 - Removing, 4-45
- Remote Procedure Call, 2-8, 6-1
- Remote Services
 - Setting Access to, 4-20
 - Setting Maximum Number of, 4-20
- Remove NFS File Access, 4-45
- REX
 - Troubleshooting, 9-69, 9-71
- RFA through NFS, B-7
- ro* Mount Options, 4-39
- RPC Services
 - /etc/inetd.conf* Entries, 4-19
 - /usr/adm/inetd.sec* Entries, 4-21
 - Activation of, 4-18
 - Security, 4-18, 4-21
- RPCGEN, 2-9
- rstatd(1M)*, 4-10
- rusersd(1M)*, 4-10
- rwalld(1M)*, 4-10

S

- Security, 4-18, 4-20—4-21
 - Considerations, 5-9
 - Limitations, 5-9
 - Master Servers, 7-16
- Servers
 - NFS, 2-1, 4-10
 - PC NFS, 4-16, 4-31
 - Preventing Access, 4-47
 - rstatd(1M)*, 4-10
 - rusersd(1M)*, 4-10
 - rwalld(1M)*, 4-10
 - sprayd(1M)*, 4-10
 - Stateless, 2-2
 - YP, 2-14
 - YP Masters, 2-16
 - YP Slaves, 2-16

setuid Mount Defaults, 4-38
Slave Servers
 Automatic Start, 7-28
 Configuration, 7-26
 Manual Start, 7-28
Soft Mounts, 4-35, 4-38, 4-40, 4-42
Software Installation, 3-6
Software Update, 4-49
sprayd(1M), 4-10
Stateless Servers, 2-2

T

Troubleshooting
 REX, 9-68—9-69, 9-71, 9-73, 9-75, 9-77, 9-79, 9-81, 9-83, 9-85, 9-87, 9-89, 9-91, 9-93, 9-95
Troubleshooting, 9-1
 /etc/passwd, 9-63
 /etc/passwd file, 9-61
 /etc/vhe_list, 9-63
 /etc/vhe_list file, 9-61
 Clusters, C-3
 Command Not Found, 9-89
 Configuration, 9-9
 Denied Access, 9-83
 Device Files, 9-95
 Errnos, 9-14
 Error Messages, 9-13
 Execution of *vhe_mounter*, 9-65
 Hardware, 9-9
 Home Node, 9-59
 Initial REX Steps, 9-69, 9-71
 Initial Steps, 9-9, 9-17
 Initial VHE Steps, 9-57
 Initial YP Steps, 9-45

Mount Fails, 9-21
Mount Point, 9-87
Named Pipes, 9-95
Network Communication, 9-10
Network Problems, 9-8
NFS, 9-11, 9-17
Performance Problems, 9-41
Permission Denied, 9-91
Programs Hang, 9-37
References, 9-6
Restricted Access, 9-33
REX, 9-81, 9-87, 9-89, 9-91, 9-93, 9-95
REX Server, 9-75, 9-81, 9-83
REX User ID, 9-77
Server Not Responding, 9-25
Text File Busy, 9-93
Unknown Host, 9-73
Unsolved Problems, 9-14
User ID Denied Access, 9-79
VHE, 9-56, 9-58—9-60, 9-62, 9-64—9-67
 vhe_mounter, 9-67
 YP, 9-11, 9-45
 YP Clients, 9-55
 ybind(1M), 9-53
 ypserv(1M) Problems, 9-51

U

/usr/adm/inetd.sec, 4-8, 4-20
UID, Setting of, 4-14
umount(1M), 4-45
Unmount File Systems, 4-45
update, 3-7, 4-49

V

VHE
 Troubleshooting, 9-57
VHE Advanced Usage
 Alternate Mount Points, 8-19
 altlogin Login, 8-17

- mouner* Login, 8-17
- Using VHE for Mail, 8-20
- VHE Configuration, 8-3
 - /etc/exports*, 8-10
 - /etc/passwd*, 8-3, 8-8
 - /etc/vhe_list*, 8-3, 8-6
 - Mounts in Background, 8-14
 - Overview, 8-2
 - Refinements, 8-14
 - Using YP, 8-3
 - vhe_mouner*, 8-11
- VHE Maintenance, 8-15
 - Adding or Deleting Nodes, 8-16
 - Unmounting File Systems, 8-15
- Virtual Home Environment
 - Advanced Usage, 8-17—8-18, 8-20
 - Advantages, 2-17
 - Concepts, 2-20—2-21
 - Configuration, 8-3
 - Disadvantages, 2-19
 - Maintenance, 8-15
 - Overview, 2-17, 8-1

W

- Wide Area Network
- Mounts, 4-39
- wsiz* Mount Defaults, 4-39

Y

- Yellow Pages, 2-11
 - Advantages, 2-12
 - Concepts, 2-13
 - Configuration, 7-14
 - Databases, 7-5

- Disabling of, 7-33
- Disadvantages, 2-12
- Domains, 2-15
- Escape Sequences, 7-7
- Log Files, 7-43
- Maps, 2-14, 7-5
- Masters, 2-16
- Slaves, 2-16
- Structure, 2-13
- Troubleshooting, 9-11, 9-45, 9-51
- Verification of, 7-32
- YP Clients, 2-14
 - Alteration of, 7-21
 - Automatic Start, 7-25
 - Configuration, 7-20
 - Manual Start, 7-25
 - Troubleshooting, 9-55
- YP Configuration, 7-14
 - /etc/hosts.equiv*, 7-11
 - Clients, 7-20
 - Clusters, C-3
 - Master Servers, 7-15
 - Propagate Maps, 7-29
 - Slave Server, 7-26
- YP Domains, 2-15
- YP Maintenance, 7-33
 - Add New Users, 7-39
 - Adding Servers, 7-39
 - Changing Master Servers, 7-40
 - Log Files, 7-43
 - Manual Modification to Maps, 7-35
 - Master Servers, 7-40
 - Modify YP Maps, 7-34
 - Non-standard Maps, 7-44
 - Password, 7-41
- YP Password, 7-41
- YP Servers, 2-14
 - Adding of, 7-39
 - Maintenance, 7-39
- ypbind(1M)*, 7-12
- ypcat(1)*, 7-12
- ypinit(1M)*, 7-12
- ypmatch(1)*, 7-13

yppasswd(1), 7-13, 7-41

yppasswdd(1M), 7-13

yppoll(1M), 7-13

yppush(1M), 7-31

ypserv(1M), 7-13

ypxfr(1M), 7-14, 7-29, 7-31

Reader Comment Card

HP 9000 Series 300 and 800 Installing and Administering NFS Services B1013-90001

We welcome your evaluation of this manual. Your comments and suggestions will help us improve our publications. Please tear this card out and mail it in. Use and attach additional pages if necessary.

Please circle the following Yes or No:

- | | | |
|--|------------------------------|-----------------------------|
| Is this manual well organized? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| Is the information technically accurate? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| Are instructions complete? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| Are concepts and wording easy to understand? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| Are examples and pictures helpful? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| Are there enough examples and pictures? | <input type="checkbox"/> Yes | <input type="checkbox"/> No |

Comments:

Date _____

Name: _____

Title: _____

Company: _____

Address: _____

City & State: _____



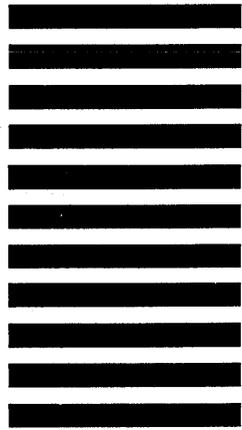


NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 37 LOVELAND, CO

POSTAGE WILL BE PAID BY ADDRESSEE

Hewlett-Packard Company
Colorado Networks Division
3404 East Harmony Road
Fort Collins, CO 80525



ATTN: Network Usability Department

Fold Here

Tape

Please do not staple

Tape



HEWLETT
PACKARD

HP Part Number
B1013-90001
Printed in U.S.A.

E0989



B1013-90001