HP-UX
System Administration Tasks

HP 9000 Series 300 Computers

HP Part Number 98594-90061

HEWLETT
PACKARD

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

Legal Notices

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, 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 direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty. A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

Restricted Rights Legend. Use, duplication or disclosure by the U.S. Government Department of Defense is subject to restrictions as set forth in paragraph (b)(3)(ii) of the Rights in Technical Data and Software clause in FAR 52.227-7013.

Use of this manual and flexible disc(s) or tape cartridge(s) supplied for this pack is restricted to this product only. Additional copies of the programs can be made for security and back-up purposes only. Resale of the programs in their present form or with alterations, is expressly prohibited.

Printing History

This manual printing date and part number indicate its current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The manual part number will change when extensive changes are made.

Manual updates may be issued between editions to correct errors or document product changes. To ensure that you receive these updates or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

September 1989 ... Edition 1. 0, Release 7.0 of the HP-UX System Administration Tasks Manual

Contents

A. Operating System Parameters

1

Introduction to System Administration

This manual explains how to perform system administration tasks. This chapter introduces you to system administration.

Using this manual assumes the following conditions:

- You are running a release of HP-UX as follows:
  - The 6.5 release. Your first task is to update HP-UX to the 7.0 release. Work through the chapter named "Updating HP-UX". If you have a release prior to 6.5, you must install the 7.0 release; you cannot update to 7.0.
  - The 7.0 release. Use the chapters according to the task you need to perform. For example, to update HP-UX, see the chapter having that name.
- Your system is configured to run the 7.0 release of HP-UX. This manual does not explain configuration. It does examine configuration in terms of meeting user needs, but you must read other material to get information about configuration.
- This manual explains fundamental aspects of administering an HP-UX system (kernel configuration, managing files, file system backup). To perform tasks related to networking, windows, applications, and languages, this manual points you to other documentation.
- The explanations of tasks assume you have the prerequisite skills (using an editor such as `vi` and executing fundamental commands such as `pwd`, `cp`, and `ps -ef`.

Finding the Chapter You Need

The chapters indicate major tasks ("Managing Groups and Users"). Task modules within chapters explain particular tasks ("Adding a User").

Ch	Major Task	Focus
2	Constructing an HP-UX System	Customizing the system so it performs according to user needs. For areas such as networking, the chapter points you to other documentation.
3	Updating HP-UX	Updating to a new release, adding software to a system, and related tasks.
4	Starting and Stopping HP-UX	System startup, attended mode, run-levels and states, shutdown, and ways to reboot a system.
5	Managing Groups and Users	Adding/deleting a group, adding/removing users, and providing customization files.
6	Managing the File System	Performing tasks related to file systems, swap space, and media (initializing media, mounting a file system, converting to long file names).
7	Managing Spooled Devices	Line printer/plotter and LP spooler configuration and management.
8	Managing Devices	Overview of managing devices. Points you to the document for installing peripherals for your series.
9	Reconfiguring the Kernel	Procedures for altering devices, swap space, and system parameters.
10	Managing an HP-UX Cluster	Making a system into a cluster, adding clients, deleting clients, accommodating context dependent files.
11	Backing Up and Restoring the System	Planning and and procedures for performing a backup or restore.

Getting Information Outside This Manual

If You Need To ...	See This Manual ...
Understand HP-UX concepts, for example, memory management or swap space	*HP-UX System Administration Concepts*
Install a new HP-UX system, or install HP-UX on an existing system	*Installing HP-UX*
Install a device or get information about a device	See the document for installing peripherals for your series. The titles can vary among series.
Use any system security features	*HP-UX System Security* manual.
Obtain reference information about HP-UX commands	*HP-UX Reference* (more than one volume exists). Most system administration commands appear in *Volume 2* and *Section 1M.*
Set up or administer networking	As appropriate, read: ■ *Using and Administering NFS Services* ■ *Installing and Maintaining NS-ARPA Services* ■ *LAN Node Manager's Handbook*
Manage X11 windows	*Using the X11 Window System*
Solve problems	*Troubleshooting HP-UX*

What Makes You a System Administrator

A **superuser** (root user) has a user identification (UID) of 0. A superuser logs in as `root` and provides the password for the root user. Saying this another way, a system administrator can log onto HP-UX as the user whose entry in `/etc/passwd` might look like this:

```
root:xE%/OqrnYf8Hg:0:1:System-Admin-Joe-Hacker-T-555-1234 :/:/bin/sh
```

Note the superuser has a user name of `root`, a UID of 0 in the third field, and a home directory of `/`, which is the root file system.

For contrast, an entry for a non-root user might look like this:

```
john:Ko%/OqrnYf8Hg:20:5:John Jones   :/users/john:/bin/ksh
```

Note the non-root user has a personalized user name, a non-zero UID (20), and a home directory in a file system under the root directory (`/users/john`).

Superuser Capabilities

The superuser (root user) can access commands and system calls not available to other users as follows:

- Execute any command in the system having *some* execution bit set in the command file's mode. A normal user can execute only those commands having the *user* or *other* execution-bit set.
- Override any protections placed on user files.
- Modify any system configuration files.
- Add or remove users to the system.
- Perform system functions, such as backups and updates.

To administer a system, you must maintain a superuser, and HP recommends using the name **root.** Be sure the superuser (root user) has a password. Share the password only with an authorized person (for example, the substitute system administrator).

The Role of a System Administrator

A system administrator serves the users in some logical group (an HP-UX cluster, a server for networking, a multi-user system having terminals, a standalone workstation).

In fulfilling this role, you provide the following types of services. In providing a particular service, you typically complete several tasks. There are no rules for this. You need to determine which tasks to perform to provide a particular service. For example, to customize HP-UX, you may edit environment files, reconfigure the kernel, and add networking services.

- Evaluate user needs and plan system resources.
- Install hardware, install or update HP-UX, and manage software (applications, tools, utilities, and libraries).
- Support local users, remote users, and clusters.
- Provide networking and communications.
- Secure a system and provide accounting.
- Backup and recover (archive and restore) files.
- Construct and customize HP-UX systems. This can include having HP-UX interact with other operating systems.
- Monitor file system use and growth.
- Detect and correct file system corruption.
- Recover from system crashes.
- Solve problems.
- Provide comprehensive expertise related to systems.

Having the Prerequisite Skills

Being an expert user of HP-UX makes system administration easier. While you need not initially be an expert, you do need a minimal level of skill to function as a system administrator.

At a minimal level, you should be able to perform the following types of tasks:

- Log in, change the password, and log out.
- Understand the file system as follows:
  - Explore directories and find files.
  - Use absolute and relative paths (for example, `/etc/motd`, `dev/dsk`, and `../lib/term`). Use commands such as `cd`, `cp`, `rm`, and `mv` to manipulate files.
  - Use commands such as `cat`, `more`, and `grep` to examine the contents of files.
- Use at least one shell (Bourne, Korn, or C). Use "dot" files to customize a shell environment (for example, `.profile`, `.rhosts`, `.x11start`).
- Use `vi` (or another editor) to edit files.

If you must function as a system administrator and you have concerns about your current level of skill, work through the HP-UX beginner's guides (HP-UX, vi, shells, X11windows, and such). The *Finding HP-UX Information: HP9000 Series 300* manual describes these guides. If you are a new system administrator and know these skills, but you lack an understanding of HP-UX, work through the *HP-UX System Administration Concepts* manual. It has fundamental information about the file system, device files, memory management, and so on.

Being Away From Your Job

There will be times when you must be away from the system you administer. During those times, you should have a substitute system administrator, a person who can perform the minimal tasks required to keep the system running.

To help the substitute, develop a customized form that explains how to handle your duties. At the least, the form needs the following types of information:

- Your name, the location of the system, the login name, and the password, which must be kept in a secure place.
- A list of duties and a schedule.
- A description of the system: CPU, memory, cards and boards, drives, printers, plotters, LAN, and so on.
- Any unusual customizations.
- Major software; especially the shell, editor, and tools.
- The names of people to contact. For example, you should list the people to call for information about such things as networking, service, and system crashes.
- You could include a description of your method of working.

Give the form to an appropriate person (possibly your manager). If the form has the password, it should be secured. If you do not supply the password, the substitute must know it. The following page contains an example, but you are encouraged to invent your own.

```
Substitute System Administrator Procedures and Information
**********************************************************

Administrator: Sam Smith,  ext 5555,  Location 1LZ9,
               Electronic mail as samh@hpfcqz, login as root,
               password is "big.one"

Substitute: Jane Jones,  ext 1111,  Location 1LZ20,  janef@hpfciq

DUTIES:
   1) Make sure system is running each morning.
   2) If system is running fine, do not do anything.
   3) Insert optical disk in optical drive at end of each
      day; backs up automatically each night; remove disk
      and store with manager each morning.  Cycle the set.
   4) In the event of a system crash, try the following:
      a) Use recovery system on a tape located with manager
         to get the system going.
      b) In an extreme case, reinstall from media kept
         with the manager.  Then, restore users from
         backup tapes.

SYSTEM INFORMATION:
   1) S 300, 350, 16 Mbytes, LAN, HP-IB at 7 and 14, SCSI, RS-232-C.
      Term is 300h, Uses Bourne shell, Printer is LaserJet II,
      Two 7937 drives, 9144 tape drive, Backup to 650A optical
      disk (manager has set of five disks).
      Server for 12 cnodes, full system with LAN.
   2) Sue and Ron run "Acme-Base".  Ed and Sally use "emacs".
      Vectras use ARPA Services.  Vectras start up in PAM.
   3) All cnodes use X11 Windows.

PEOPLE TO CONTACT:
   1) Manager is John Baker, ext 4444
   2) Maintenance is Kathy Grey, ext 4321
   3) If you have a crash, call Marv Hayes, ext 1234
   4) Each day, visit with cnodes about any problems.
      If not serious, list them and give to me later.
      If serious, accommodate the problem according to
      the tasks manual.
```

Using the System Administration Manager (SAM)

Release 7.0 of the HP-UX system provides a System Administration Manager (SAM). SAM replaces `/etc/reconfig` (which existed on releases prior to 7.0) and provides the following benefits:

- Instead of executing commands from a shell, you work through menus that guide task selection and facilitate data entry.
- Tasks are easier to perform because you need not remember or type complex commands.
- You get a rich set of functions, and those functions provide significant options, control, and power.
- You get a complete set of Help Screens. You can get context-sensitive help from every location in SAM.
- You can use SAM on Series 300 *or* Series 800 systems without relearning anything. The menus are the same. Data-entry screens fit the system.
- The `/etc/newconfig/Update_info/reconfig_sam` file has information you should read.

This module explains how to use SAM. Run SAM as you work through the material. Explore the things being explained.

When you execute `sam` to start up the System Administration Manager, you will see that performing a task amounts to moving down through menu options until you reach a data-entry screen. There, you enter the data required to perform a task.

This module does not describe everything SAM helps you do. Subsequent chapters explain system administration tasks. If SAM can do a task, the SAM procedure appears first; the manual procedure appears second.

The last part of the module has a tutorial session. Although brief, the tutorial shows how to use major features of SAM: the menu-oriented task selection; the help screens; and the data-entry screens.

Running the System Administration Manager

After installing the 7.0 release (or after updating your system from release 6.5 to 7.0), you can use SAM. (You are encouraged to run SAM while working through this section to see how it works.)

Become the Superuser to Run SAM

1. Log in as `root`. You typically do this in run-level 2 (or perhaps 3). If SAM requires being in the single-user run level, it will get into that run level or tell you to do so.
2. To run SAM, execute:

   ```
   sam
   ```

The Initial Display Has the Following Parts

```
SAM                  System Administration Manager

Highlight an item and then press "Return" or "Select Item".

  ...

<Other items and information appear here>
```

If you do not see this screen (menu), SAM did not start up, and you will need to use your problem-solving skills to determine why.

Do This Exploration and Then Continue

Examine the screen. Note the parts, especially the options (for example, `How To Use SAM`). Take time to explore, noting all parts of the screen display. Use Tab or arrow keys to highlight an option. Press Return to initiate actions. Use softkeys according to their function. Basically, SAM provides menus that lead to a screen for entering task-related data.

The exploration and Help Screens may provide all the information you need to use SAM. If they do not, the remaining sections provide specific information.

Selecting a Menu Item

SAM works by having you select a menu option related to the task you want to perform. An arrow at the end of an option means you get another menu.

1. Press (Tab) (or an arrow key) enough times to highlight your option.
2. Then, press (Return) or the softkey for `Select Item` (the (f4) key).

From the Main Menu, you can perform tasks related to the following menu options, which appear near the center of the window.

```
Users ->
Groups ->
Auditing and Security (Trusted System) ->

File Systems ->
Peripheral Devices ->

Networks/Communications ->
Cluster Configuration ->
Kernel Configuration ->

How To Use SAM
```

Now is a good time to select the `How To Use SAM` option. By reading the information in the screen that appears, you may not need to read the information shown later in this section.

Types of SAM Screens

SAM has four types of screens.

Type of Screen	How It Works and How You Respond
Menu Screens	Show options (tasks) you can select. Use (Tab) or arrow keys to highlight an option and press (Return) (or (Select Item)). An arrow following an option, ->, means the option leads to another Menu Screen. An ellipsis following an option, ... , means you get a Data Entry Screen. Otherwise, an option displays information or performs a task without interaction.
Data Entry Screens	Provide a form that has fields for entering data. Fill in or modify the fields, using (Tab), (Return), or arrow keys to move among fields. Press `Perform Task` when you finish entering data.
Help Screens	Show information about an item. Each highlighted item has a Help Screen. Some help screens list valid entries for a field. Press `Select Item` to place a valid entry in the field. In some Help Screens, underlined or highlighted items provide "hypertext" help when you press (Return) (or `Select Item`). Press `Exit Help` to return to SAM.
Feedback Screens (Messages)	Show messages from SAM as follows: ■ Error messages appear when you enter invalid data. They tell you which field to correct. ■ Progress messages indicate that SAM is performing a task. SAM does not execute all tasks immediately. Wait until a message disappears. ■ Confirmation messages let you enter `y` to confirm taking an action or `n` to cancel it. ■ Termination status messages indicate a task has completed.

Types of Data Fields

When you reach a data-entry screen, you see questions or prompts, each followed by a line (the data field). Enter data in each field required for SAM to complete the task. For optional fields, you can enter data, but you need not do so.

Type of Field	How It Works and How You Respond
Normal Data-Entry Fields	A line under the field indicates its length. Enter the data and use (Tab) or the arrow keys to move to the next field. After you enter data in required fields (and desired optional fields), press `Perform Task` to have SAM perform the task or press `Exit Task` to cancel the data you entered.
Shiftable Fields	Some fields can have data that extends beyond the underline. The field shifts as required to let you enter more data. When this happens, a < or > indicates more data lies to the left or right of the visible field. Use the left and right arrow keys to move to the right or left in these fields.
Scrollable Fields	Some forms and help screens contain more information than can be shown at one time. For example, some help screens list the possible entries for a given field. SAM shows five items at a time. Use the (Prev), (Next), (Tab), or arrow keys to scroll through the fields (lists).
Pop-up Fields	When the data for a field dictates what appears in a data-entry screen or when you provide more than one piece of information for a given field, you get a pop-up field. Press `Done` when you finish entering data in the pop-up field.
Display-only Fields	You get displayed information that you cannot modify.

Escaping to a Shell

When you run SAM, and then need to work in a shell, press the softkey for `Shell`. You leave SAM and enter a shell environment (Bourne, Korn, or C). (SAM does not work with job control.)

Getting the shell prompt (usually **$** or **%**) confirms being in the shell. Execute commands as required. Type `exit` to return to SAM at the point where you escaped.

Navigating within SAM

To navigate SAM, begin at the Main Menu and select options from Menu Screens until you reach a Data Entry Screen. Enter data in the fields of the Data Entry Screen. Use (Tab) (or (Shift)-(Tab)) to move among fields. When you finish entering data, press the `Perform Task` softkey to have SAM perform the task.

You can use keys and softkeys to navigate.

- Softkeys initiate the actions indicated by their corresponding labels.
- Keys initiate the actions indicated by the (Keycap).

The action taken by a softkey may duplicate the action taken by a key. For example, pressing the softkey for `Select Item` takes the same action as pressing (Return). Beyond this, the items in the following tables show specific ways to navigate SAM.

Navigating in Menus

Typing This ...	Takes This Action ...
Space Bar	Highlights the next item.
Tab Key	Highlights the next item.
Shift-tab	Highlights the previous item.
Arrow Key	Highlights an adjacent item.
Return Key	Selects the current item.
Type a minimal unique prefix	Selects the item. Does not work when the two menu items have the same first name. Also, within Item Selection Help, typing a minimal prefix may find an item, but it does not select an item.

Navigating in Data-Entry Screens

Typing This ...	Takes This Action ...
Space Bar	Makes a space.
Tab Key	Moves the cursor to the next field, does not scroll.
Shift-tab	Moves the cursor to the previous field, does not scroll.
Arrow Keys	Move the cursor to the nearest field in the indicated direction. Arrow keys can also shift and scroll. They keys facilitate navigation better than Tab in shiftable and scrollable fields.

The table continues on the next page.

Typing This ...	Takes This Action ...
Home Key	Moves cursor to first field in the data-entry Screen
Prev Key Next Key	Scrolls a scrollable array as indicated by the key label.
Insert Char Key	Toggles insert mode.
Delete Line Key	Clears the current field.
Delete Char Key Backspace Key	Deletes the current character.
Return Key	Works as follows: ■ Moves cursor to the first field on a next line, if there is a next line. ■ If there is no next line, the Return key is equivalent to "Done". ■ In a scrollable array or field, the Return key scrolls.

Navigating With a Mouse

In an X11 hpterm window, you can activate a softkey action by clicking on the corresponding softkey label.

Refreshing the Screen

From any location in SAM, Ctrl-L refreshes the screen.

The Major Structure of SAM

To help you see possibilities for navigation, the chart on the next page shows the first two or three levels of SAM, beginning with the Main Menu. Besides suggesting possibilities, the chart shows you what SAM does.

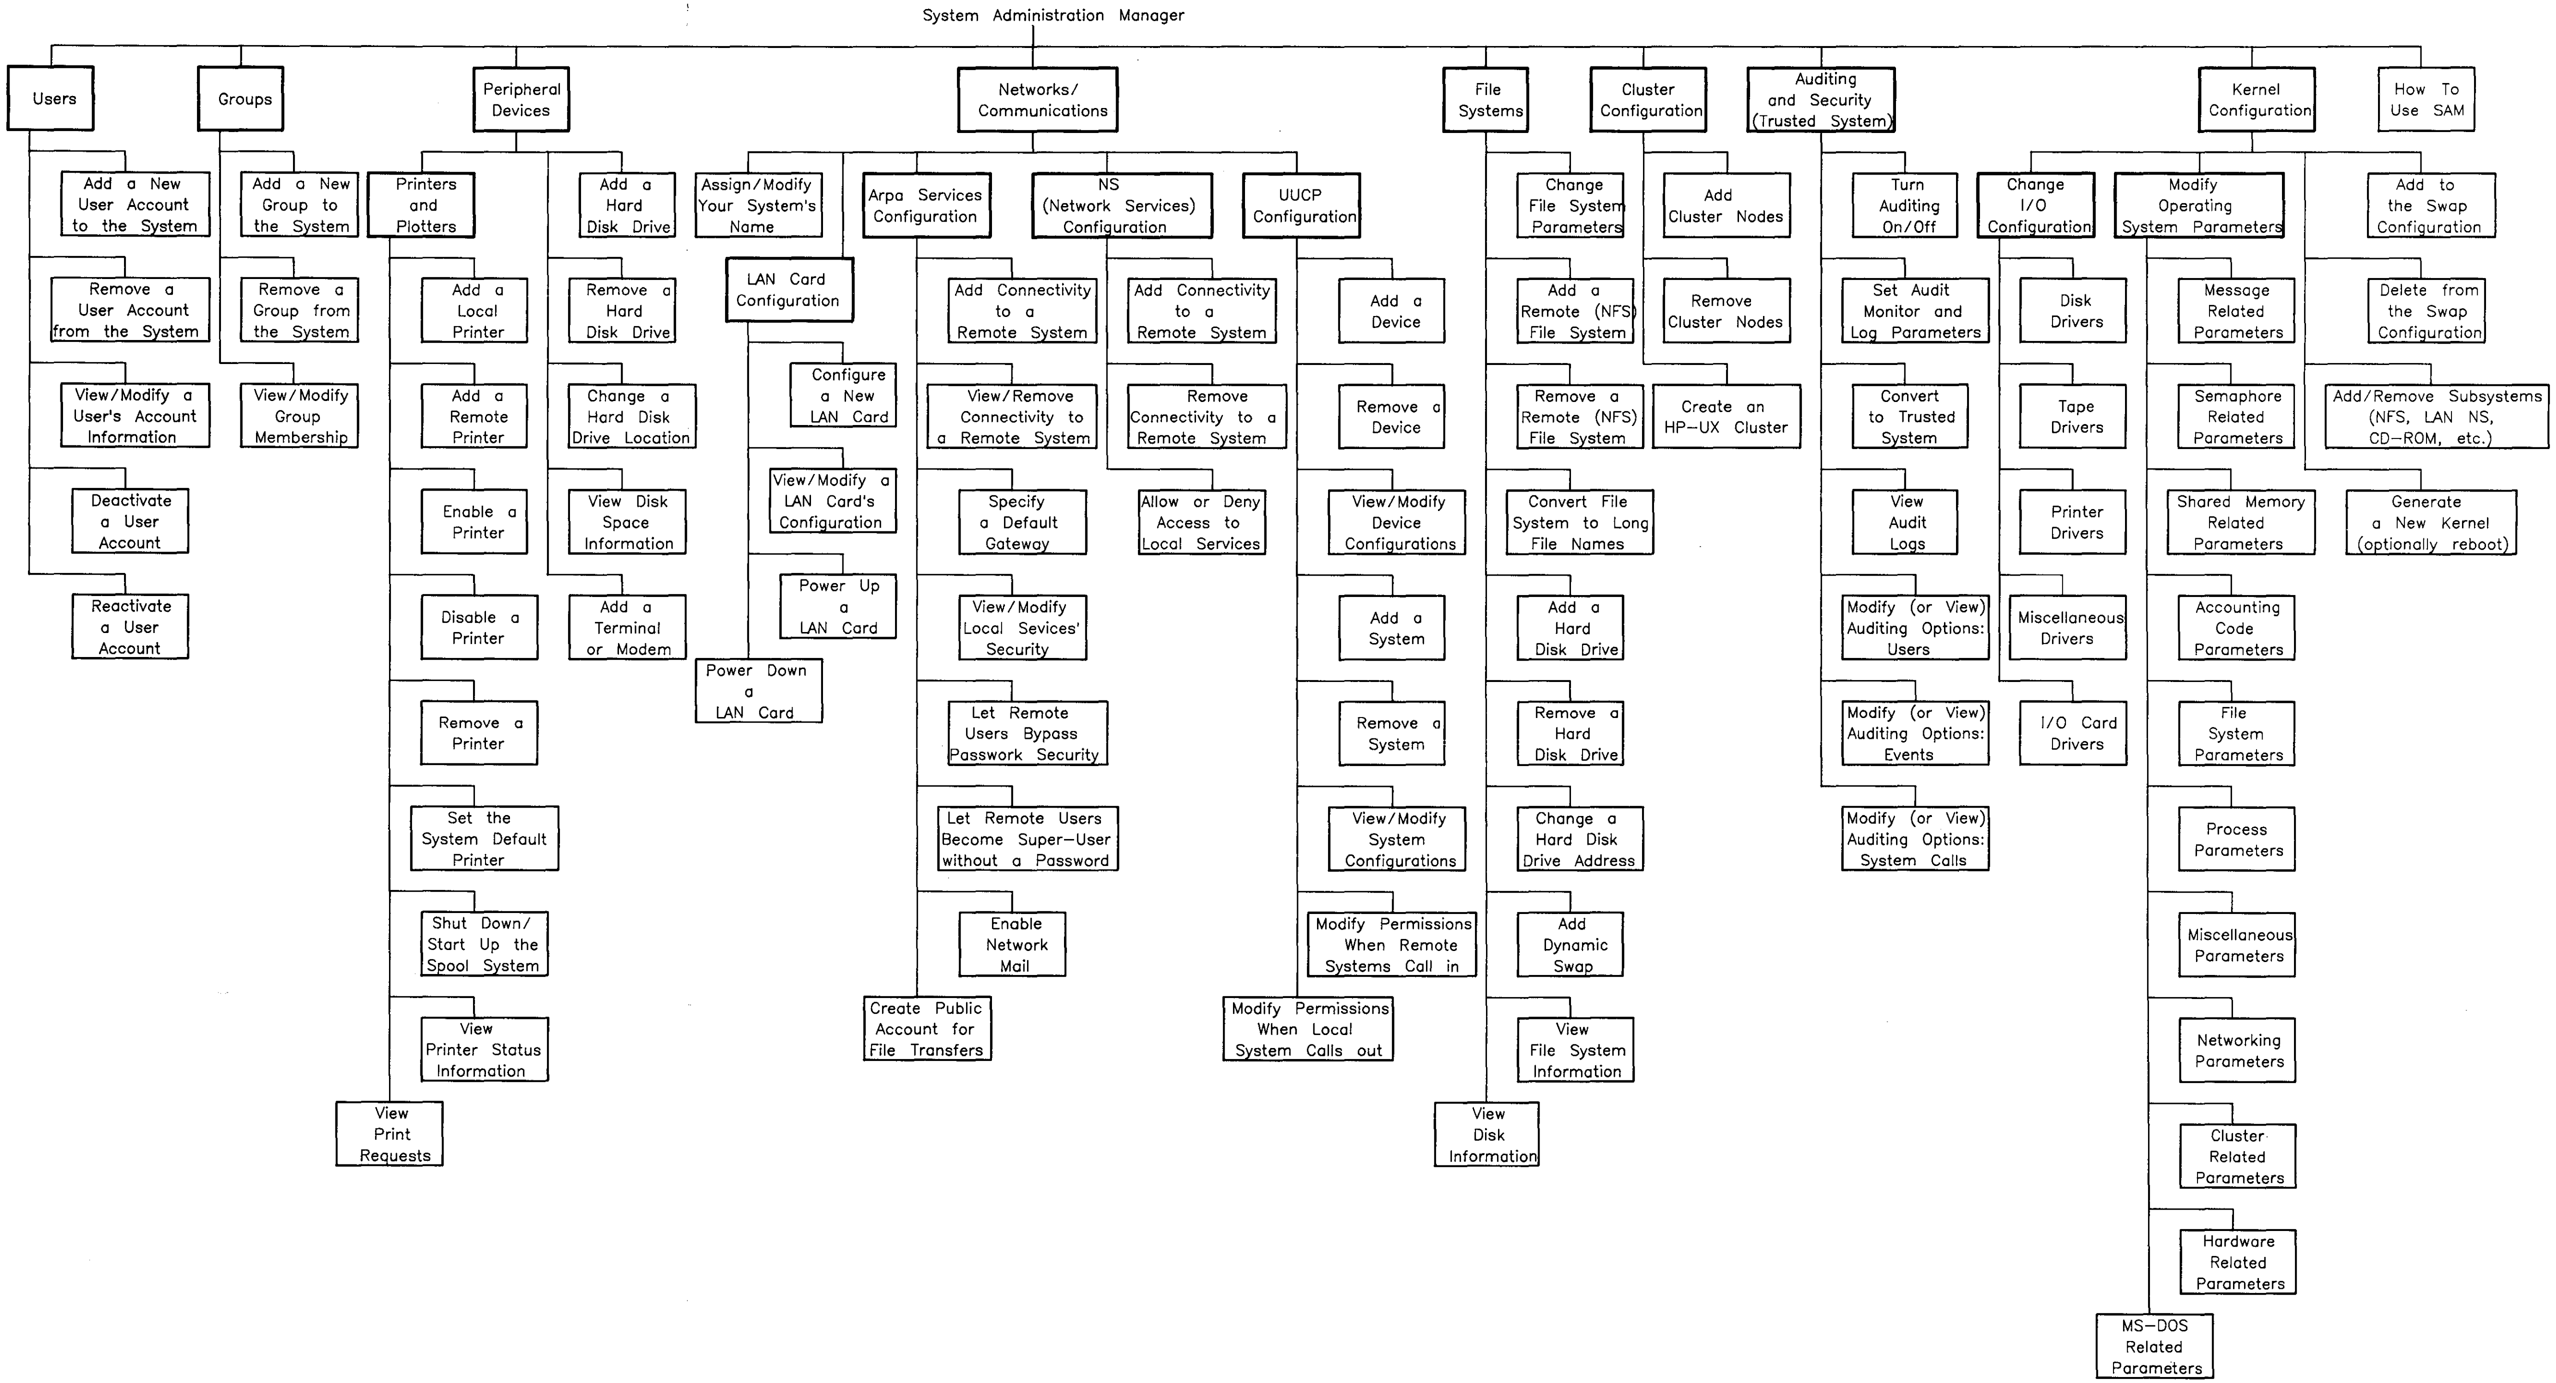
System Administration Manager
Users
Add a New User Account to the System
Remove a User Account from the System
View/Modify a User's Account Information
Deactivate a User Account
Reactivate a User Account
Groups
Add a New Group to the System
Remove a Group from the System
View/Modify Group Membership
Peripheral Devices
Printers and Plotters
Add a Local Printer
Add a Remote Printer
Enable a Printer
Disable a Printer
Remove a Printer
Set the System Default Printer
Shut Down/Start Up the Spool System
View Printer Status Information
View Print Requests
Add a Hard Disk Drive
Remove a Hard Disk Drive
Change a Hard Disk Drive Location
View Disk Space Information
Add a Terminal or Modem
Networks/Communications
Assign/Modify Your System's Name
LAN Card Configuration
Configure a New LAN Card
View/Modify a LAN Card's Configuration
Power Up a LAN Card
Power Down a LAN Card
Arpa Services Configuration
Add Connectivity to a Remote System
View/Remove Connectivity to a Remote System
Specify a Default Gateway
View/Modify Local Sevices' Security
Let Remote Users Bypass Password Security
Let Remote Users Become Super-User without a Password
Enable Network Mail
Create Public Account for File Transfers
NS (Network Services) Configuration
Add Connectivity to a Remote System
Remove Connectivity to a Remote System
Allow or Deny Access to Local Services
UUCP Configuration
Add a Device
Remove a Device
View/Modify Device Configurations
Add a System
Remove a System
View/Modify System Configurations
Modify Permissions When Remote Systems Call in
Modify Permissions When Local System Calls out
File Systems
Change File System Parameters
Add a Remote (NFS) File System
Remove a Remote (NFS) File System
Convert File System to Long File Names
Add a Hard Disk Drive
Remove a Hard Disk Drive
Change a Hard Disk Drive Address
Add Dynamic Swap
View File System Information
View Disk Information
Cluster Configuration
Add Cluster Nodes
Remove Cluster Nodes
Create an HP-UX Cluster
Auditing and Security (Trusted System)
Turn Auditing On/Off
Set Audit Monitor and Log Parameters
Convert to Trusted System
View Audit Logs
Modify (or View) Auditing Options: Users
Modify (or View) Auditing Options: Events
Modify (or View) Auditing Options: System Calls
Kernel Configuration
Change I/O Configuration
Disk Drivers
Tape Drivers
Printer Drivers
Miscellaneous Drivers
I/O Card Drivers
Modify Operating System Parameters
Message Related Parameters
Semaphore Related Parameters
Shared Memory Related Parameters
Accounting Code Parameters
File System Parameters
Process Parameters
Miscellaneous Parameters
Networking Parameters
Cluster Related Parameters
Hardware Related Parameters
MS-DOS Related Parameters
Add to the Swap Configuration
Delete from the Swap Configuration
Add/Remove Subsystems (NFS, LAN NS, CD-ROM, etc.)
Generate a New Kernel (optionally reboot)
How To Use SAM

Tips for Using SAM

If You Get Lost in the Structure

`Main Menu` Gets back to the top-level menu.

`Exit Task` Returns to the previous screen, canceling any changes you made to the current screen.

If You Need Help

When you do not know how to do something, press the `Help` softkey. SAM knows its context and provides help screens for options and data-entry fields.

If You Get an Error Message

Error messages can appear at the bottom of menus and data entry screens. You must clear them before continuing. You will see the statement:

```
Press space bar to continue.
```

Pressing the space bar returns SAM to a normal state.

An internal error can display one of the following messages:

```
SAM internal error occurred.  Please call your HP representative.
Internal error 512
```

Beyond this, status messages appear in the lower middle portion of the screen. Pressing a space or alpha key continues.

If You Lose Power to the Terminal

SAM may not function normally when you power up the terminal again. Proceed as follows:

1. Type your interrupt key, probably CTRL-C

   If this does not work, try Shift-Del

2. Answer y to the question:

   ```
   Do you really want to quit?
   ```

3. You get a shell prompt. Execute:

   ```
   reset
   ```

   You may need to use Ctrl-J.

4. Restart SAM by executing:

   ```
   sam
   ```

An Illustrative Example of Using SAM

You can practice using SAM by working through this tutorial.

Prerequisites and Conditions

- Log in as the root user.
- Start up SAM by executing `sam`. You will see the main screen, which is a menu.

Adding a User to the System (Example)

Work through the tasks in the following sections (possibly 3-8 minutes).

Select the Appropriate Menu Item

1. In the Main Menu, the item `Users ->` is the one to select. It is already highlighted, so press (Return), or the softkey for `Select Item` (the (f4) key). The menu screen for `Users` appears:

   ```
   Add a New User Account to the System ...
   Remove a User Account from the System ...
   View/Modify a User's Account Information ...

   Deactivate a User Account ...
   Reactivate a User Account ...
   ```

2. While you are in this menu screen, take a look at the help screens for these menu items by highlighting a menu item and pressing the `Help` softkey. Each menu item has its own help screen.
3. Continuing the task of adding a user, use the arrow keys or the (Tab) key to highlight `Add a New User Account to the System` and press (Return), or the `Select Item` softkey. The data entry screen appears.

Complete the Data Entry Screen

1. You complete nine fields. You must enter data in each field unless it is marked `(optional)`. Each field has a help screen.
2. In the first field ("Login name"), type `utest`.
3. Use (Tab) (or arrow keys) to move the cursor to the next field ("Primary group name"). Whenever possible, SAM provides a default value for fields, as it has done here. Here, the default is "users", and for convenience, the tutorial uses it. For now, move the cursor to the next field ("Home Directory").
4. When you entered the "Login name" in the first field, SAM added that to the `/users` directory. Consequently, the default home directory for this new user is `/users/utest`. We'll use this default, so move the cursor to the next field ("Start-up program").

The steps continue on the next page.

5. You have pre-defined choices for the start-up program (or shell). Press the `Help` softkey to see the list. The default start-up program is the Bourne shell (`/bin/sh`) Instead of using it, give the user the Korn shell by highlighting `/bin/ksh` and pressing (Return) or the `Select Item` softkey.
6. You should now be in the field labeled "Login with X11 windows? (y or n)". Use the default value for this field.
7. From this point on in the screen, the fields are optional, which means you you can leave them blank.

Take a Moment to Compare Values

Before you continue, compare your screen values with the table values.

In the Field ...	**You Should Have ...**
Login name	`utest`
Group name	`users`
Home directory	`/users/utest`
Start-up program	`/bin/ksh`
Login with X11 windows? (y or n)	`n`
Real name	blank
Office location	blank
Office phone	blank
Home phone	blank

Perform the Task

1. Press the `Perform Task` softkey. A pop-up screen appears asking you to enter a password for the new user.
2. Enter the password `test1` and press (Return) or the `Done` softkey. Notice that the password does not appear as you type it.
3. After you enter the password, SAM asks you to reenter the password. Once again, type `test1` and press (Return) or the `Done` softkey.
4. A progress message appears, letting you know that SAM is adding the user. When SAM has completed the task, another message appears telling you that `utest has been added to the system`. Press the space bar to make this final message disappear.
5. You are now back in the data-entry screen. You can add another user, but do not do so now. Instead, press the `Exit Task` key to return to the menu screen for `Users`.

Verify the Results

1. In the `Users` menu, highlight the menu item
   `View/Modify a User's Account Information`
   and press (Return) or `Select Item`.
2. A pop-up screen asks you to enter the login name of the user to view or modify. Type `utest` and press (Return) or `Done`. Incidentally, pressing the `Help` key displays a list of the current users from which you can select the one you want to view or modify.
3. The account information of the user you just added to the system (`utest`) appears in the data-entry screen. SAM has successfully added the user.

Your Turn ...

Using SAM, remove the user you just added. If you need help, the procedure for doing this is on the next page.

How to Remove the User You Added

1. From the Main Menu of SAM, use (Tab) or the arrow keys to highlight the menu item `Users`, then press (Return) or `Select Item`.
2. From the `Users` menu, use (Tab) to highlight `Remove a User Account from the System` and press (Return).
3. A pop-up window appears, asking for the login name of the user to remove. Enter the login name (`utest`) and press (Return) or `Done`.
4. SAM asks if you want to remove all files and directories belonging to this user from the system. The default is `y` (yes). You do want to remove all of this user's files and directories, so do not change the default. This is the only entry you can make in this screen.
5. Press the `Perform Task` softkey. A confirmation message appears to ensure this is really what you want to do. It is, so press `y` (yes) to continue with the process.
6. A progress message appears, letting you know that SAM is removing the user. When SAM has completed the task, another message appears telling you that `utest has been removed from the system`. Press the space bar to make this message disappear.
7. Pressing `Exit Window` exits this data-entry screen. Then, press the same key (now labeled `Previous Menu` to move to the Main Menu.
8. To exit SAM, press the `Exit SAM` softkey, respond with a `y` to the confirmation message, and you are out of SAM and back in your shell.

A Strategy for Using SAM

At this point, you know how to use SAM. You also have some information about what SAM does. With this in mind, the following items suggest a strategy for using SAM:

- HP recommends using SAM to administer a system. To this end, the modules present the SAM procedure first.
- SAM does not accommodate every task you need to perform (for example, backup a file or check file system consistency). To accommodate this, the modules always contain a manual procedure for performing a task.
- While performing a task with SAM, you may encounter a situation that SAM cannot accommodate. If this happens, you can escape to a shell and perform the task (or part of the task) manually. For example, on a Series 800 computer, SAM can add a printer, but it does not make the device file for the printer. On a Series 300 computer, SAM does make the device file.
- If you escape to a shell to perform a task and do not know how to proceed, you can get information by reading the manual procedure that follows the SAM procedure.
- SAM attempts to accommodate your needs, but you must realize that administering a system can require problem solving to account for situations. You may find that SAM does not always perform tasks according to your needs, but that should not constitute a rationale to stop using SAM.
- You should be able to use SAM efficiently and use manual procedures in the following situations:
    - SAM cannot perform a task.
    - You know (as an expert) how you want to customize a functionality.
    - You encounter a corner-case that requires escaping to a shell to perform a task manually.

2

Constructing an HP-UX System

This chapter explains how to construct and customize an HP-UX system. Your work can affect the root directory, `/`, as follows:

- Adding software can add commands to `/bin`, `/usr/bin`, and other "bin" directories.
- Adding software can add files to `/lib`.
- Making device files alters `/dev`. The chapter named "Managing Devices" has basic information about device files.
- You make many customizations by editing files in `/etc` and `/usr`.
- You typically add users to `/users`, but you have alternatives, such as creating a file system named `/other` and adding users there.
- During the construction process, you may create special-purpose file systems (for example, `/net` or `/archive`).

Besides file systems, the root directory contains essential files (for example, `/hp-ux`, `/SYSBCKUP`, and `/SYSDEBUG`). Reconfiguring the kernel changes these files. Your customizations may require reconfiguring the kernel.

The file systems provide shells, utilities, tools, languages, applications, and environments. You can work in the Bourne, Korn, or C shells; use the `vi` or `emacs` editors; and use various languages (for example, C, HP-BASIC, Pascal). HP-UX supports graphics, windows, networking, and system security, among other things. You have much freedom in doing the construction.

Using the Appropriate System Run Level

This module discusses the system run levels (or states) because, later, as you perform tasks, you may need to get into a certain run level. (The chapter on "Starting and Stopping HP-UX" has a complete description of run levels and system states.)

Prerequisites and Conditions

- To do system administration and still accommodate users, you set the system run level.
- Most of the time, the system functions at run-level 2 (the multi-user state). To work on the file system, you often bring the system to run-level *s* (the administrative or single-user state).
- Run levels range from 0 to 6. You most often operate the system in run-level 2 (you can also use run-level 3). States other than 1, 2, or 3 exist for special purposes.
- Do not run the system at run-level 0. It is reserved for system installation.
- For a client in an HP-UX cluster, changing the run level has no effect on other clients in the cluster.
- Changing the run level of the cluster server of an HP-UX cluster can affect all active clients. In particular, changing the run level to *s* makes the server unable to respond to requests from clients. Thus, the clients panic and halt.
- Before you change to run-level *s* to do system administration tasks, tell each user when the system will go down.
- It is often better to work during the off hours to perform tasks that require run-level *s*.

The Single-user Run Level

Run-level *s* (the single-user or administrative state) provides a run level for doing tasks that require a quiet system (for example, no open files). Each time you perform a task, note if it requires being at run-level *s*.

Getting into the Single-user or Administrative Run Level

1. Do *not* use `init s`.
2. Use `shutdown` instead.

The directions often recommend using `shutdown 30`, which provides a 30-second grace period.

Getting out of the Single-user Run Level

After you do something in the single-user state, you typically reboot the system, run a file system check, or perform another task. This action has two implications:

1. You may run one or more followup commands that also require being in the single-user state. Here are some examples:

`reboot`	*Starts up the system again, typically returning you to the multi-user state.*
`reboot -h`	*Brings the system down to a complete halt.*
`fsck`	*Performs a file system consistency check.*
`swapon`	*Activates the swap space.*
`mount -a`	*Mounts all file system in /etc/checklist.*

2. Eventually, the system must return to the appropriate run level. This typically means you execute `reboot`.

Making a Recovery System

After you install HP-UX and before you consider user needs or make customizations, consider ways to recover from a crash or a forgotten root password. A recovery system lets you start up a subset of HP-UX when you cannot start up the normal HP-UX. From this limited system, you can attempt to repair the total system. While this is an optional task, consider its importance.

Prerequisites and Conditions

- You need not get into the single-user state.
- Do not use flexible disks. Use a 150-foot cartridge tape. A 600-foot tape works, but it takes much longer to create the system.
- Each time you update the operating system, create a new recovery system.
- On a secured system, lock the recovery system tape. Otherwise, starting up from the recovery system gives you an administrative superuser mode with no auditing.
- The root block-device file should be one of:

`/dev/dsk/0s0`	*The preferred device file name.*
`/dev/hd`	*A holdover from earlier times.*
`/dev/root`	*Used by some system administrators.*

  The tape drive character-device file should be one of:

`/dev/update.src`	*Usually exists on the system*
`/dev/rct`	*Used by many system administrators*
`/dev/rct/c0`	*Less common but acceptable*

- Your system must have the *KERN_BLD* fileset. Listing `/etc/filesets` should show it. If the fileset is not present, use the `/etc/update` utility to add it.

Manual Method for Making a Recovery System

1. Login as the root user.
2. Insert the cartridge tape in the drive and wait for it to load.
3. Execute `/etc/mkrs` according to the following syntax:

    `mkrs` [-*v*] [`-f` *rcdev*] [`-r` *rootdev*] [-*m series*]

    For example, the following command uses the indicated tape drive device file, disk drive device file, and verbose mode.

    ```
    mkrs -v -f /dev/update.src -r /dev/dsk/0s0
    ```

    By default, executing the following command uses `/dev/update.src` and `/dev/dsk/0s0` without verbose mode.

    ```
    mkrs
    ```

4. The following things happen, or could happen:
    a. The process takes 1-2 hours on a 150-foot tape, up to 6 hours on a 600-foot tape.
    b. The examples did not use the *m series* option (for example, `m` 300). The command normally knows your system. If it does not, you get an error message. Run the command again and include the option.
    c. If you get error messages related to "command not found" or "missing device files", update your system to include the filesets or make the device files. Then, try again.
5. When the process completes, remove the tape, label it, and store it in a safe and secure place. Later, test the recovery system to see if it works. To do that, you need to be in the single-user state. See the later module named "Recovering from a System Crash".

Recovering from a System Crash

If HP-UX gets into a non-functional state, you can try to restart the system or repair it by running `/etc/fsck`. If these attempts fail, and you cannot get the system to run, you can use your recovery system. If you do not have one, you may lose data because you need to reinstall HP-UX from your original tape.

Prerequisites and Conditions

- Turn every device OFF.
- Use the recovery system for your release of HP-UX. Do not use a recovery system for the 6.5 (or earlier) release to recover a system for the 7.0 release.
- On the cartridge tape that contains your recovery system, have the arrow point away from *SAFE*.
- Turn on the cartridge tape device used to create the recovery system. At the least, use a device having a character device file at the same bus address.
- Turn on devices that have swap space on your original system. If your system has several mass-storage devices, and they contain swap space, turn them on as well.
- Insert the cartridge tape having the recovery system in the tape drive and wait for the busy light to remain off.

Procedure for Using a Recovery System

1. Turn the computer on and immediately hold down the space bar until you see `keyboard` in the list of items appearing down the left side of the display. This puts you in **attended mode,** which shows available systems down the right side of the display.
2. Wait until all systems appear (1-2 minutes). You might see:

   ```
   7937: 1400, 0, 0
   1H SYSHPUX
   1B SYSBCKUP
   9144: 0700, 0, 0
   2H SYSHPUX
   ```

3. Whatever you see, select the system for the device having the cartridge tape. (For the example, type 2H and do not press Return.)
4. Wait for the startup process to complete (about 10 minutes). The recovery system does not need certain files, but you see the messages anyway. For example, you may see:

   ```
   prod#.... .... .ignored
   ```

5. Wait for the `fsck` to complete. The recovery system does a file system consistency check of the root device. This takes about 20 minutes.
6. You eventually see a recovery menu. Follow its directions. You can call a help facility if you need directions for doing a task.
7. When you finish the recovery process (or the recovery cannot be completed), shut down and halt the recovery system by selecting the following option from the recovery menu:

   ```
   Exit Recovery System and Reboot root file system
   ```

8. Finally, remove the tape and, again, store it in a safe and secure place.

Continuing When the Recovery System Does Not Work

If your system crashes and neither a file system consistency check nor using the recovery system fixes the problem, you have two choices:

1. In the menu for the recovery system, you have the following option:

   ```
   Work in a shell to perform manual recovery
   ```

   While HP does not recommend using this method, you can elect to use it. Try this method only when the recovery process did not succeed and you want to try this option before doing anything else. If you use this option, the following section has some information.

2. Your other choice is to reinstall the system from your original media. This means you begin anew. After you install the system, you need to copy backed up file systems to the new system. You may need to perform additional construction and customization.

Using the "Work in a shell ... " Option for Recovery

1. Use your recovery system to start up a subset of HP-UX. (The earlier section named "Procedure for Using a Recovery System" explained this.)
2. Select the option named:

   ```
   Work in a shell to perform manual recovery
   ```

3. When you get into the shell, the following stepped sections suggest what to do. This manual assumes you have the expertise to perform the tasks.

Step 0: Assumptions for Using the Option

You can use this option to recover from specific, localized, problems on your file system.

The procedures in subsequent steps make the following assumptions:

- You cannot boot your regular system; you suspected a problem; and you used the recovery system unsuccessfully.
- Your root device is called `/dev/real.root` (block device file) and `/dev/rreal.root` (character device file).
- Your recovery device has two names: `/dev/root` is the block device file for the recovery device and `/dev/rroot` is the character device file for the recovery device.
- Use `ls -l /dev` to determine what device files are actually present on the recovery system. Subsequently, you should make any adjustments in names.
- Root device, in the following procedure, refers to the device that is root under normal circumstances (the hard disk drive associated with your root file system).
- Automatic recovery attempts all the following steps, but one or more steps may fail for reasons that you can correct if you have sufficient knowledge of HP-UX.

Step 1. Check the Critical System Files

To get your root volume to a bootable state while still running on the recovery system, check the files described in the next paragraphs.

/bin/sh. Copy the version of this file on the recovery system to the root volume and relink `bin/rsh` by executing:

```
cp /bin/sh /disc/bin/sh
ln /disc/bin/sh /disc/bin/rsh
```

/etc/init. Copy the version of the this file on the recovery system to the root volume by executing:

```
cp /etc/init /disc/etc/init
```

/etc/inittab. If `inittab` is corrupted, `init` might fail. Save the `inittab` file for later editing and then create a single line `inittab` by executing:

```
mv /disc/etc/inittab /disc/etc/inittab.save
echo "is:s:initdefault:" > /disc/etc/inittab
```

Type the second line exactly as shown (including the quotes).

In an HP-UX cluster this file is a CDF. It must be referenced as:

`/disc/etc/inittab+/`*rootserver_name*.

/etc/ioctl.syscon. If you have changed the device used as the console, `/etc/ioctl.syscon` can be incorrect or corrupt. Remove the file because the system will create a correct file during the next startup procedure.

```
unlink /disc/etc/ioctl.syscon
```

In an HP-UX cluster this file is a CDF, so you reference it as:

`/disc/etc/ioctl.syscon+/`*rootserver_name*

Continue on the next page.

/dev/console. This file is linked to `dev/syscon` and `dev/systty`. It can be corrupted, resulting in an unbootable system.

In an HP-UX cluster `/dev` is a CDF, so reference it as:

```
/disc/dev+/localroot/console
/disc/dev+/localroot/syscon
```

and

```
/disc/dev+/localroot/systty
```

Disregarding the group name and date, you should get the following line when listing the file:

```
crw--w--w- 3  root 0 0x000000  /disc/dev/console
```

If you do not get the line, re-create the files as follows:

1. Remove `console`, `syscon`, and `systty` by executing:

   ```
   rm /disc/dev/console /disc/dev/syscon /disc/dev/systty
   ```

2. Use `mknod` to recreate `/dev/console` as follows:

   ```
   /etc/mknod /disc/dev/console c 0 0x000000
   ```

3. Link the files by executing:

   ```
   ln /disc/dev/console /disc/dev/syscon
   ln /disc/dev/console /disc/dev/systty
   ```

Continue on the next page.

/hp-ux. If the kernel file is corrupted, boot from a backup kernel *or* copy `hp-ux.min` from the recovery system to the root device.

In an HP-UX cluster this file is a CDF, so reference it as:

`/disc/hp-ux+/`*rootserver_name*

If you have a backup kernel (`SYSBCKUP`) on your root device and if the backup kernel is not corrupted, you can boot your system from `SYSBCKUP` as follows:

1. Reboot your system, holding down the space bar until you see the available systems.
2. Select the `SYSBCKUP` option (typically `2H`).
3. When you have rebooted, create a new kernel by using `/etc/config`) *or* by copying the backup kernel to `/hp-ux`.

If you have no other usable backup kernel, try to boot by using `/hp-ux.min` from the recovery system. You can copy `/hp-ux.min` to the root disk by executing:

```
cp /hp-ux.min /disc/hp-ux
```

Then, try to boot from the root disk.

Continue on the next page.

Step 2. Fix Other Problems

If you have rebooted from the root volume, you may need to fix some problems. You may not have any of these problems, but you do need to check them. If they exist, use the corrective procedures.

Corrupted inittab. If `/etc/inittab` was corrupted, edit the version that was saved. Once edited, move the file back to `/etc/inittab`. Until you test the new `inittab`, make the default state be run-level *s*. Later, you can switch to other run levels by using `init` *x* where *x* equals a run level from 0 to 6.

Lost Files. If you lost other system files, use `/etc/update` to add the lost files to the system again. Running `update` requires the following files:

```
/bin/cat                 /bin cnodes         /bin/getcontext
/bin/echo                /bin/find           /bin/grep
/bin/sh                  /bin/mkdir          /bin/make
/bin/pwd                 /bin/cpio           /bin/ps
/usr/bin/lifcp           /usr/bin/tcio       /bin/rm
/etc/update              /etc/sysrm          /etc/config
/etc/devnm               /etc/mount          /etc/reboot
/etc/setmnt              /usr/bin/lifrm      /usr/bin/madecdf
/etc/regen               /system/TOOL/lif_or_not
/system/TOOL/rebuild.ckerns
```

Check your system to ensure these files are present.

Then, you can get missing files from your recovery tape by executing the following sequence of commands:

```
mkdir /disc              only if /disc does not already exist
mount /dev/root /disc
cp /disc/bin/mkdir /disc/bin/pwd /disc/bin/cpio /bin
cp /disc/usr/bin/lifcp /disc/usr/bin/tcio /usr/bin
cp /etc/update /etc/sysrm /etc
```

Boot Area Corruption. The boot area of the root device may be corrupted. If you suspect this, create a raw device file for root (assuming none exists) and copy the boot area from the recovery system to the root device by executing:

```
dd if=/dev/rfd of=/dev/real.root count=1 bs=8k
```

Having Backup Commands. The recovery tape has the commands normally used for backup (`cpio`, `tcio`, and `find`). If your root disk does not have these commands, and you want to restore from backups, mount the recovery tape and copy the commands to the root disk.

Step 3. Return to Recovery Tool

Type `exit` or `^D` to return to the recovery tool. Then reboot the system using the `Exit Recovery System and Reboot root file system` option. Remove the recovery system tape.

Using Alternatives for Recovering a System

Your system crashed, you tried your recovery system, you tried the manual method offered in the recovery system menu, and; you still cannot get the system up and running. This section describes alternative procedures for recovering from a crash.

If you have backup media for `/users` and other files that are not on the tapes you purchased, you can use the original tape for release 7.0 of HP-UX to restore your system. Proceed as follows:

1. Reinstall the system, using the *Installing HP-UX* manual.
2. Customize the system again, using the appropriate tasks described in this manual (for example, adding users).
3. Restore `/users` and other files contained on your backup tapes. Use the methods described in the chapter named "Backing Up and Restoring the System".

This process can help you get going again.

Evaluating User Needs and Configuring a System

You need to conceptualize the existing situation and evaluate user needs. Then, you develop ways to accommodate those needs. The following items suggest possibilities:

1. If you have no hardware or software, you plan the system and then purchase the hardware and software.

   Before you purchase anything, survey what your users need to function effectively.

2. If someone purchased the components of a system, you install and customize the system.

   Since nothing has been installed, ask people what they want and then customize the system as best you can. If the purchased configuration is not correct, you may need to buy additional (or alternate) devices, cables, and cards.

3. If you inherited an existing system, it is obviously working, but it may have shortcomings you want to correct.

   Ask people what they need, customize the system as best you can, and plan ahead for needed hardware and software.

Providing Features the User Wants

Users typically discuss features or entities. They want multiple windows (a feature) or the `Acme Word Processor` (an entity). Users seldom discuss kernel parameters or device files. With this in mind:

1. Survey the users' wants.
2. Translate the wants into what HP-UX provides. For example, having users say they want the DOS operating system means you install the SoftPC or DOS Coprocessor products.
3. Continue this process until you have enough information to configure a system.

Planning for Required Resources

Translate the information gathered from users into a system (for example, an HP-UX cluster with networking and X11 windows). Then, obtain the resources (hardware, software, budget, and expertise). The following examples illustrate this:

1. Having a multi-tasking system that gives users the equivalent of standalone workstations requires setting up an HP-UX cluster.
2. Accommodating users who run data-entry applications requires setting up a multi-user system that has terminals.

Before you can configure a system, you need to translate user needs (and wants) into a system (a cluster, a networking server, a multi-user system, a workstation, and so on). Having a system in mind helps you do the specific configuration.

Configuring an HP-UX System

During configuration, you translate components such as computers, monitors, and printers into exact products, options, and bundles. The table shows the idea.

If You Planned For ...	Your Configuration Might Require ...
A computer	HP 9000 Model 350MH 98589E
Mass Storage Devices	HP 7936H 307 Mbyte fixed (CS80) HP 9122C 3.5 inch flexible (CS/80)
A Graphics Device	HP 7575A w/17570A (HP-IB)

You typically configure the following types of systems.

Type of System	Description
Standalone Workstation	A system that contains an SPU, monitor, keyboard, mouse, mass-storage device, printer, and other devices according to the specific needs of the user. The user often administers the system.
HP-UX Cluster	A cluster server supports clients by providing resources. Resources include printers, plotters, modems, software environments, and so on. More than being terminals, the clients function as standalone workstations.
Multi-user System with Terminals	A system, much like a standalone system, but it supports terminals via RS-232 ports.
Networking Server	A system that is specifically configured to provide networking service.

Getting Information about Configuration

Configuring a system requires an exact specification of products as replacement for components. Take your time. Seek expert help. Do not install any hardware until you believe the configured system will work. It can become difficult (or impossible) to install an incorrectly configured system.

The following document has information about functionality, products, enhancements, languages, integration, database management, communications, and documentation.

HP 9000 HP-UX Operating System: A Technical Supplement.

The following data sheet contains information about SPUs, monitors, keyboards, and operating systems.

HP 9000 Series 300 Hardware Technical Data

The following booklet explains configuration steps; discusses minimum and recommended systems; suggests accessories and peripherals; provides ordering information; and contains tips. The section on networking, for example, contains tables that relate Network Name, Product Number, Product Description, Supported Transports, Functional Description, and Configuration Requirements.

HP 9000 Series 300 Hardware Configuration Guide

Installing and Testing Hardware

After you accommodate user needs by planning and configuring their systems, set up the computer. Verify what it contains because you may have a bundled package. Especially, verify the cards in the computer. Then, install all peripherals, devices, and additional cards. Begin by assembling the following items:

1. The System Processor Units (or SPUs), devices, accessories, and cards.
2. The documentation for every SPU, device, accessory, card. Check this because some components have loose papers that get misplaced (for example, the *Computer Information* card for the SPU has information such as the serial number and lanic ID).
3. The manual for installing peripherals for your series.

Install the SPUs, devices, accessories, and cards according to their documentation. Hardware documentation should explain installation and testing. The following table shows the idea for some typical devices.

To Install This Device ...	**Use This Documentation ...**
Keyboard	*HP-HIL Keyboard User's Guide*
Mouse	*HP Mouse Owner's Guide*
ID Module	*HP 46084A and 46084-69901 ID Module*
Human Interface Board	*Configuration Note: Human Interface Board*
Any peripheral	The documentation for that peripheral and the installation guide for peripherals for your series.

Test the SPU, devices, accessories, and cards according to their documentation. Do not install or update any software until you know the hardware works.

If you have problems, contact your HP representative.

Installing or Updating a Release of HP-UX

After you have installed and tested appropriate hardware, you can install the HP-UX operating system. If you are running the 6.5 release of HP-UX, you can update to the 7.0 release.

Installing the 7.0 Release of HP-UX

An installation loads files from a cartridge tape (or other supported media) onto a mass-storage device, constructs the HP-UX file system, builds an initial HP-UX kernel, and provides an initial login. You need not be a system administrator to perform this task. You do not need to be running any release of HP-UX. In contrast, you can be running any release of HP-UX including the current one.

Updating to the 7.0 Release of HP-UX

An update to the 7.0 release loads products from a cartridge tape (or other supported device), incorporates the filesets into the file system, and reconstructs the existing HP-UX kernel. You must be the system administrator and be running the the 6.5 release. After the update, the `/etc/newconfig` and `/etc/newconfig/Update_info` directories contain important files.

Continuing According to Your Situation

Decide whether you need to install or update HP-UX. Then, continue according to the following items:

- To install HP-UX, work through the *Installing HP-UX* manual. Then, return to this chapter and construct your system.
- To update HP-UX, work through the later chapter named "Updating HP-UX". Then, return to this chapter and perform those tasks that were affected by the update. For example, after an update, you may need to provide system security or edit some environment files.

Setting the System Clock

Constructing an HP-UX system requires setting the system clock. Many commands use the clock to accomplish their tasks.

Prerequisites and Conditions

- Monitor the system clock daily, at the least. Reset it if it is off by more than a minute.
- Keep the system clock set to the time and date broadcast by a station that accurately gives the Greenwich time (for example, WWV in the United States).
- You need not reset the system clock on a Series 300 just because you power down the system. Series 300 computers have a battery that keeps the clock current.
- On an HP-UX cluster, setting the system clock is the same as with other HP-UX systems. However, the system clock is cluster-wide. Clocks for clients synchronize with the cluster server's clock as they join the cluster, and they remain synchronized.

Setting the Time Zone

Only the superuser can change the system clock.

Set the time zone environment variable (TZ) before you set the current time and date in the environment. Set the time-zone value with a variable declaration (as shown later) in three possible files: `/etc/csh.login`, `/etc/rc`, and `/etc/profile`. The time zone, TZ, can also be set from an application program with the `tzset` library routine.

In these files, the format looks like this:

```
TZ=XXXHYYY          use with /etc/rc
TZ=XXXHYYY          use with /etc/profile
export TZ
setenv TZ XXXHYYY   use with /etc/csh.login
```

where:

XXX An alphabetic abbreviation of the standard time zone, usually three letters in length. For example, `MST` for Mountain Standard Time.

H The difference between standard local time and Greenwich Mean Time, in hours. Fraction hours indicate minutes (for example, 3:30 for Newfoundland). Positive hours move West from Greenwich (for example, use 7 for Mountain Standard Time). Negative hours move East from Greenwich (for example, -9:30 for South Australia).

YYY An alphabetic abbreviation of the daylight time zone for your area, usually three letters in length (for example, `MDT` for Mountain Daylight Time). Delete this part if Daylight Savings Time is not observed in your geographic area.

Within the United States, the following examples show the idea:

- In Eastern time zone, use `TZ=EST5EDT`
- In Central time zone, use `TZ=CST6CDT`
- In Arizona, where Daylight Savings Time is not observed, use `TZ=MST7`

For more information on setting the time zone environment variable, refer to `TZ` under the *environ(3)* entry in the *HP-UX Reference* manual.

CST6CDT now has two different meanings because the United States makes the transition to Daylight Saving Time on the first Sunday in April and Canada makes the same transition on the last Sunday in April. The `/usr/lib/tztab` file has this information but it requires `CST6CDT#Canada` to distinguish between the two.

Continue on the next page to set the time and date.

Setting the Time and Date

Once you set the time zone variable, terminate the `cron` process (if it is running) and execute the `date` command.

1. Kill the `cron` process.

   On an HP-UX cluster, terminate `cron` on each clients. To terminate `cron`, locate the `cron` process information by executing:

   ```
   ps -ef | grep cron
   ```

   This identifies the Process ID (PID) for `cron`. To determine all `cron` processes for all clients in an HP-UX cluster, execute:

   ```
   cps -ef | grep cron
   ```

   With this information, terminate `cron` by executing:

   `kill` *pid*

   where *pid* is the process ID associated with `cron` (for example, `16442`).

2. Set the correct time and date (using the `date` command) by executing:

   ```
   date MMddhhmm{yy}
   ```

   where:

   a. `MM` is a two-digit integer representing the month. For example, 03 represents March.

   b. `dd` is a two-digit integer representing the day of the month. For example, 02 represents the second day of the month.

   c. `hh` is a two-digit integer specifying the current hour in terms of a twenty-four hour clock. For example, 03 specifies 3:00 am and 14 specifies 2:00 pm.

   d. `mm` is a two-digit integer specifying the number of minutes past the stated hour. For example, 04 specifies four minutes past the hour.

e. `{yy}` is an optional two-digit integer specifying the last two digits of the current year; this parameter may be omitted if the year is already correct. For example, 87 specifies 1987 as the current year.

When `date` executes, it shows the time and date on your screen.

3. Restart `cron` if you terminated it in step 1.

   To restart `cron`, execute:

   ```
   /etc/cron
   ```

Possible Problems When Changing the System Clock

The `make` program is sensitive to a file's time and date information and to the current value of the system clock. While setting the clock forward will not affect `make`, *setting the clock backward by even a small amount may cause* `make` *to exhibit extremely bizarre behavior.* Avoid setting times earlier than the current system clock's value.

Making incremental backups depends heavily on the correctness of the date because incremental backups occur in relation to a dated file.

Altering the system clock can cause unexpected results for routines scheduled by `cron`. When setting time back, `cron` doesn't run until the clock "catches up" to the point from which it was set back. For example, if you set the clock back from 8:00 to 7:30 (which is *not* advised), `cron` will not begin executing until the clock again reads 8:00. If you set the clock ahead, `cron` attempts to "catch up" by immediately executing all routines scheduled to run between the old time and the new time. For example, if you set the clock ahead from 9:00 to 10:00, `cron` immediately executes all routines scheduled to run between 9:00 and 10:00.

Customizing Your System

After you install (or update) HP-UX, you then complete several tasks as the root user to make the system function (add users, create groups, edit environment files, and so on). This module describes the **environment files.** Basically, your system reads the environment files during startup and system operations to determine how it should function. Thus, the environment files determine your shell, password, security, electronic mail, editor, paths to directories having commands, and so on.

Prerequisites and Conditions

- Your system may not start up properly if you add incorrect entries to `/etc/inittab`, `/etc/rc`, or `/etc/passwd`. In addition, your system needs certain files to operate properly. While you edit the environment files, work very carefully and do not remove or alter files you may think are unnecessary.
- Do not modify the structure of the Context Dependent Files (CDFs) in an HP-UX cluster. If you copy directories, use **`find`** with the **`-hidden`**, **`-depth`**, and `-print` options to "catch" CDFs.
- If you update HP-UX (as opposed to install), the process does not overwrite some customized files. Instead, it puts them in `/etc/newconfig`. The `/etc/newconfig/README` file has information.

  1. Compare files in the new `/etc` with your customized files in `/etc/newconfig` to identify files to edit.
  2. Edit the new or customized files as required. Customizing the environment files assumes you understand editing and shell programming.
  3. Move customized files you want to keep from `/etc/newconfig` into `/etc`.

The Major Environment Files You Can Edit

The next sections describe the environment files. You can edit any of them manually. If you use SAM to perform a high-level task, you may be asked to supply information that relates to an environment file. When you ask SAM to perform the high-level task, SAM uses the information to edit the related environment files for you. For example, if you use SAM to add a user, SAM edits the `/etc/passwd` and `/etc/group` files.

/etc/inittab

This text file contains entries for the run levels supplied or created on your system when you start up HP-UX. The `/etc/init` command provides process control initialization in accordance with entries in `inittab`. Editing this file controls how your system runs because `init` lives from system startup until you shut down HP-UX or have a system crash. The "System Startup and Shutdown" chapter has more information.

When you add a new remote terminal to your system (for example, a fourth terminal), add a `getty` entry to `/etc/inittab` that has the following form:

```
04:2:respawn:/etc/getty tty04 H #terminal at persons desk
```

When you start up the system, the remote terminals receive a `login:` prompt. If you use SAM to add a terminal, the process adds the `getty` line to `/etc/inittab`.

In an HP-UX cluster, the file must be a CDF that contains one subfile for each client. (The "Managing HP-UX Clusters" chapter has more information.)

/etc/rc

This shell script defines several actions taken during the startup procedure. To see all the actions, examine the script, noting the parts and comments. For most systems, you should define the system *hostname*. The "Starting and Stopping HP-UX" chapter has more information. If you add networking, you may need to edit `/etc/rc` according to the documentation for networking.

/etc/passwd

This text file controls who can log into your system. The file typically contains lines such as the following:

```
root:xE5/0qrnYf8Hg:0:1:SYSTEM-ADMIN-Sam-Fix-T-555-1234      :/:/bin/sh
uucp:BjYSZ8Mem.QrE:5:5:hp-ux-2-hp-ux:/usr/spool/uucppublic:/usr/lib/uucp/uucico
nuucp::6:1:0000-uucp(0000)         :/usr/spool/uucppublic:/usr/lib/uucp/uucico
uuhpfclp:wMjyzp7jSVLGw:5:5:uu-frm-lp:/usr/spool/uucppublic:/usr/lib/uucp/uucico
acme:DkciXk2h5TOLw:1002:5:mp-login:/usr/spool/uucppublic:/usr/lib/uucp/uucico
ace:6WDckG/5q1deQ:1001:5:mc-login:/usr/spool/uucppublic:/usr/lib/uucp/uucico
lan:pTCiSbhh7cyTE:21:1:Network-Access:/:/bin/rsh
group:*:200:20:-------------------------------------:/users/guest:/bin/rsh
joe:Ko4uDcqeJ1IHs:205:20:Joe Cool                 :/users/joe:/bin/sh
archive:OrBr6eUXqJZw.:209:20:Rob Keeper           :/users/archive:/bin/ksh
jane:,..:210:20:Jane Doe                          :/users/jane:/bin/csh
```

Every user on the system (local or remote) must have an entry in this file. How a line gets added to the file depends on the method used to add the user (SAM or manual). Colons separate fields in a line. From left-to-right, the line contains the following information:

user name — A name such as `joeh`, `ktm`, `sue-ann`.

password — Up to 9 characters, at least one of which must be a numeric or special character.

*ID*s — Two values that identify the user and his or her group, (for example, `200:10` for user 200 in group 10).

real user name — A proper name such as `Susan Jones`.

home directory — Typically `/users` for people and something else for applications (for example, `/usr/spool/uucppublic` for `uucp`).

execution command — Usually a shell command such as `/bin/ksh` (Korn shell) and something else for applications.

The chapter named "Managing Groups and Users" has more information.

/etc/group

This text file identifies the users that form a group, associates group IDs (GIDs) with group names, lists users, and associates those users with a group name and a group ID. Parts of the file might have lines as follows:

```
root::0:root
other::1:root,daemon,uucp,who,date,games,sync
bin::2:root,bin,daemon,lp
 ...
mail::6:root
 ...
users::20:joe,jane,sue,bill
```

The chapter named "Managing Groups and Users" has more information.

/etc/motd

This text file contains the "mesage-of-the-day". The message appears each time a user logs in if the user's personal customization file (`/etc/profile` for Bourne and Korn shell users or `/etc/csh.login` file for C shell users) has the following line:

```
cat /etc/motd        # message of the day
```

All users see the message in an HP-UX cluster because `/etc/motd` is not a Context Dependent File (CDF).

Edit the file as often as necessary. The example shows the idea.

```
Monthend reports due this week.
Going away party for Leslie on Wednesday.
Sign your time-sheets by Friday.
```

/usr/news

The root user owns this directory, which ships empty. The system administrator can (optionally) edit files and place them in the directory. Proceed as follows:

1. Edit a news file and place it in `/usr/news`. Say anything you wish.
2. Make sure the `/etc/profile` file for Bourne and Korn shell users has the following lines:

   ```
   if [ -f /usr/bin/news ]
   then news -n    #notify if news.
   fi
   ```

   Users having the the `/etc/csh.login` file need the following lines:

   ```
   if ( -f /usr/bin/news ) then
      news -n    #notify if new news.
   endif
   ```

3. After logging in, the user sees the following message:

   `news`: *news_filename*

4. Instruct users to type `news` to see the message. They can also type `news` and the name of the file they wish to view (the directory can contain more than one file).

The system administrator can change the permissions for the directory to allow any user to broadcast a message. The procedure remains the same. The *news*(1) command has options that control how often a user receives notification. See the *HP-UX Reference* manual.

/etc/profile or /etc/csh.login

These shell scripts execute automatically during the login process. The `/etc/profile` file executes for Bourne, Korn, and restricted shell users. The `/etc/csh.login` file executes for C shell users.

While the two files have differences related to differences among shell programming commands and syntax, you can customize such things as the path for executing commands, time zone, terminal type, and so on. Study the files and read their comments to see the possibilities.

There are other files, `/etc/d.profile` and `/etc/d.cshrc`, that people can use to make their personal customizations. A later section discusses those files.

/etc/wtmp

The system uses this binary file to keep a history of logins, logouts, and date changes. The system automatically creates this file, which grows without bound. Check the file regularly and empty it. Use `/etc/last` to access the contents of the file. For example, you might get the following message for the root:

```
$ last -1 root
root     console    Wed May 24 12:59   still logged in
```

For the console, you might get messages like the following ones:

```
$ last console
LOGIN    console    Wed May 24 12:59 - 12:50 (00:00)
 ...
root     console    Fri Feb 24 14:37 - 11:07 (66+19:29)
```

(The *wtmp*(4) and *last*(3) entries in the *HP-UX Reference* manual have more information.)

This file must be a CDF for HP-UX clusters (`/etc/wtmp+`). The CDF must contain a subfile for each client. (The *cdf*(4) entry in the *HP-UX Reference* manual has more information.)

/etc/btmp

If this binary file exists, the system uses it to keep track of *bad* login attempts. To get the information, you must create this file by executing:

```
touch /etc/btmp
```

The file grows without bound, so check and empty it regularly. Use `lastb` to access the contents of the file. See the previous section for examples. The *lastb*(1M) and *utmp(4)* entries in the *HP-UX Reference* manual have more information.

This file is a CDF for HP-UX clusters (`/etc/btmp+`). The CDF must contain a subfile for each client.

/etc/utmp

Do not remove this file.

The `who` command uses this file, which is created by the system. The file contains a list of current users and system startup information.

The *utmp(4)* entry in the *HP-UX Reference* has more information.

This file is a CDF for HP-UX clusters (`/etc/utmp+`); the CDF must contain a subfile for each client.

/etc/securetty

If it exists, this text file specifies the `tty` files on which the `root` user can log in. You must explicitly create this file and place the `tty` device file names in it to use this feature.

In an HP-UX cluster, create this file so it is a CDF (`/etc/securetty+`). The CDF must contain a subfile for each client. The entries in each client's subfile may differ.

/usr/lib/terminfo

This subsystem identifies terminal capabilities for programs such as the `vi` text editor. The subsystem defines terminal attributes for all Series 300, Series 800, and HP-supported terminals.

The subsystem also contains terminal attributes for terminals *not* supported by Series 300 HP-UX, but you use them without HP support.

/etc/checklist

This text file contains a list of mountable file systems and swapping devices. When you do not specify a device file for an `fsck` command, the command performs its checks on the HFS file systems listed in `etc/checklist`.

The `sam`, `diskusg`, `mount`, `umount`, `swapon`, and `fsclean` commands also use `/etc/checklist`. If you use SAM to add a disk or NFS file system, add the entry to `/etc/checklist`.

The `/etc/checklist` file ships with a single device file name, `/dev/dsk/0s0`, which corresponds to the hard disk on which you installed the root file system. If you create additional file systems and you want to mount them during system startup time, add entries for each additional disk drive containing a file system.

If you have an HP-UX cluster, this file must be a CDF (`/etc/checklist+`). If you are a client in an HP-UX cluster, the file should be empty unless you have local swapping on your cnode.

The "Managing the File System" chapter has more information.

Some $HOME Files

Users may request help in putting the following files into their home (login) directories.

.profile
: executes each time the user successfully logs in using the Bourne shell, Korn shell, or restricted shell. *A Beginner's Guide to Using Shells* has examples.

.kshrc
: If present, this Korn shell script supplements actions taken by the `.profile` file. Use the script to define the local shell variables, commands, aliases, and file permissions you want to apply to every spawned Korn shell. (*A Beginner's Guide to Using Shells* has examples.)

  For HP-UX to find this file, define and export the *ENV* environment variable in the `.profile` file.

```
# Set up Korn shell initialization with one-shot magic

ENV='${START[ (_$- = 1) + (_ = 0) - (_$- != _${-&%&%*i*}) ]}'
START=~/.kshrc
export ENV
```

.cshrc
: executes when a new C shell starts, and a new shell starts when a user logs in or issues a shell escape. The *Beginner's Guide to Using Shells* has examples.

.login
: executes when a C shell user logs in, but after `.cshrc` executes. (The *Beginner's Guide to Using Shells* has examples.)

.environ
: executes when a user logs in using PAM.

Examples of `.profile`, `.login`, and `.environ` ship under the names `etc/d.profile`, and so on. If you use SAM to add a user , the process places the appropriate environment file (or files) in the user's home directory.

$HOME/.exrc

This text file maps terminal characteristics and sets up new key definitions so you can use features such as arrow keys with the `ex` family of HP-UX editors (`vi`, `ex`, and so on). If it exists, `.exrc` must in the user's home directory. The editor searches for the file and uses the definitions to create extra editor features.

The `.exrc` file functions only when you do not define the *EXINIT* in the `/etc/profile` or `$HOME/.profile` files.

You get an example of the file in `/etc/d.exrc`. You may want to customize the file and provide it to new users by default. An appendix to "The vi Editor" article in the *HP-UX Concepts and Tutorials* manual has more information.

/etc/man_fix

This new utility lets you manage the man pages. For example, you can remove, compress, or reformat them. To some extent, this utility replaces the `/etc/catman` utility, which is described next. To get more information, see *man_fix*(1M) in the *HP-UX Reference* manual.

/etc/catman

Executing `catman -z` (`-z` for uncompressed) creates the cat files for the online reference manual pages. You get improved performance and sacrifice some disk space. The *catman*(1M) entry in the *HP-UX Reference* manual has more information.

You have the following alternatives for creating online documentation:

1. Create the processed manual pages by executing `/etc/catman` with no parameters. The process can take five or six hours to complete; so you might want to run it in the background at night.
2. Create selected sections of the processed manual pages by executing `/etc/catman` *sections* (where *sections* is one or more logical sections in the *HP-UX Reference* manual, for example: `1`, `9`, or `1M`.
3. Do not execute `/etc/catman`. If you create the `/usr/man/cat` directories, the first execution of `man` for a manual entry processes the entry, adds to the appropriate `cat`, and uses that file in subsequent accesses.

   The following script creates the `cat` directories:

   ```
   cd /usr/man
   for num in 1 1m 2 3 4 5 7 8 9
   do
    mkdir cat$num
   done
   ```

HP recommends the third alternative if you can spare some disk space but do not want to use any more than is necessary. With this "build-as-you-go" alternative, the system fills the `cat` directories as users access them with `man`.

If you used the first or second alternatives, so the man pages exist, you could remove the `nroff` source files if you need to recover some of the disk space required for the formatted version of the manual.

/usr/lib/tztab

Use the `/usr/lib/tztab` file to handle changes to and from summer time zones (Daylight Savings Time) in the United States and to accommodate future changes to these adjustments. This file is a "Time Zone Table" that contains the value of the TZ variable described above, followed by lines detailing transitions in the time zone adjustment.

See the *tztab*(5) entry to get details about modifying this file.

/etc/ttytype

The `tset` command uses this file as a database of terminal types on your system. Edit this file when you add types of terminals and modems to your HP-UX system. Edit the lines so they reflect the terminal types attached to your system. For example, you might have:

```
300h console
2397 tty00
2397 tty01
```

This file is a CDF in HP-UX clusters (`/etc/ttytype+`).

/etc/issue

This file contains information printed by a terminal's `getty` process prior to the login prompt. Seeing the message during the login process helps users identify the current system.

Edit `/etc/issue` so it has a line something like the following one. Adjust the value to reflect your release.

```
Welcome to Release 7.0 of HP-UX
```

This file is a CDF in HP-UX clusters (`/etc/issue+`).

Configuring the X11 Window System

As installed (or updated), HP-UX contains the files required to use the X11 window system. As shipped, you get a root window and one working window; so you will probably want to do some customization.

Prerequisites and Conditions

- Your system must have a bit-mapped monitor. See the *Using the X11 Window System* manual for other system requirements.
- At present, SAM does not perform major tasks related to X11 windowing. In the `Add a User` option, you can set up a user so the windowing system starts up automatically when a user logs in.

Manual Procedures

You can configure the X11 environment as extensively as you like. Do any of the following things:

- Work through the *Beginner's Guide to the X Window System.* That document describes most customization you want to do.
- Work through the *Using the X11 Window System* manual. That document has most of the information you need. When it does not (for example, programming), the manual points you to other documentation.

You need not do all the customizing at once. You can change the windowing environment over time, and users can do much of this work. You may need only to act as a resource.

Networking Systems

Networking lets a local system interact with remote systems. The networking can provide services that enhance the ways users can function. You can choose among many networking services. You need to decide which services you want.

Hardware Prerequisites and Conditions

- You must have appropriate hardware. The information provided by HP in material such as the *HP 9000 Series 300 Hardware Configuration Guide*) provides information about the following things:
  - Local Area Networking Products
  - Wide-Area Networking Products
  - Networking to Digital Equipment Corp (DEC ™) Systems
  - HP Shared Resource Management (SRM)
  - Networking to International Business Machines (IBM) Systems

  This information lets you see what is available.
- The documentation for each networking service describes the hardware requirements. Check these requirements against the system you purchased. Some bundled systems may already have the required hardware. You may need to install additional hardware. Pay particular attention to licensing requirements.
- Accommodate the hardware requirements and install required products before you install, update, or configure any software.

Software Prerequisites and Conditions

You typically encounter three situations:

1. The filesets loaded during installation include software for UNIX ™ Asynchronous Communications (`cu`, `uucp`, and `uux`). You need not purchase additional software, but you do need to customize the system according to information found in the *UUCP Concepts and Tutorials* manual. That manual has all the information you need.
2. You may have the filesets required to configure some networking services. The filesets required for NS-ARPA/Berkley, and NFS services are frequently bundled with a system. If the filesets were not bundled with the system, you need to purchase them and update your system.
3. You probably do not have the filesets required for the NS-VAX services, for example. Should you need these services, obtain the software and update your system.

The "Series 300 Networking" section in the *HP 9000 Series 300 Hardware Configuration Guide* has information as follows:

Item	Description
Network name	For example, ARPA-Berkley, RJE
Software products	For example, 50952CL, 50967A
Supported transports	For example, IEEE 802.3, V.35 Bisync
Services provided by the product	For example, `ftp`, a file transfer utility.

Whatever you decide to use for networking, the product has its own documentation for installation, customization, and programming. For example, to use the NS-ARPA Services, see the *Installing and Maintaining NS-ARPA Services* manual. The *Finding HP-UX Information HP9000 Series 300* manual lists all documentation related to networking.

Procedures for Setting Up Networking

The documents mentioned in earlier sections describe the manual methods. Besides using any manual methods, SAM can perform several networking tasks.

LAN Options Provided by SAM

System Administration Manager
↓
Networks/Communications ->
↓
LAN Card Configuration

You can select and then perform the following tasks:

- Configure a New LAN Card
- View/Modify a LAN Card's Configuration
- Power Up a LAN Card
- Power Down a LAN Card

ARPA Options Provided by SAM

`System Administration Manager`

↓

`Networks/Communications ->`

↓

`Arpa Services Configuration ...`

You can select and then perform the following tasks:

- `Add Connectivity to a Remote System`
- `View/Remove Connectivity to a Remote System`
- `Specify the Default Gateway`
- `View/Modify Local Services' Security`
- `Let Remote Users Bypass Password Security`
- `Let Remote Users Become Super-User without a Password`
- `Disable or Restart Network Mail`
- `Create Public Account for File Transfers`

NS Options Provided by SAM

`System Administration Manager`

↓

`Networks/Communications ->`

↓

`NS (Network Services) Configuration ...`

You can select and then perform the following tasks:

- `Assign/Modify Your System's Nodename`
- `Add Connectivity to a Remote System`
- `Remove Connectivity to a Remote System`
- `Allow or Deny Access to Local Services`

Installing, Configuring, and Using Communications

On an HP-UX system, the terms **communications** and **networking** can mean the same thing, or they can refer to different means of having systems interact.

In this manual, networking refers to services provided by networking products (for example, ARPA-Berkley Services). Communications refers to HP-UX commands that help you communicate with other people. The following table shows the communication commands.

Command	How It Helps You Communicate
`/usr/bin/uucp`	Provides file transfer among UNIX ™ systems, which includes HP-UX. Additional commands within the `uucp` facility let you send messages. The facility uses RS-232C.
`/usr/bin/mailx` `/usr/bin/mail`	Provides electronic mail capability among users having NS-ARPA Services. The `mailx` and `mail` commands are user agents of the `sendmail` and `uucp` facilities. An alternative here is to use the `elm` facility.
`/usr/bin/news`	A command and directory for helping users stay abreast of announcements.
`/etc/wall` `/etc/cwall`	Used by the system administrator to warn users about sudden changes in the run level of the system, use `cwall` on an HP-UX cluster.
`/bin/write`	Copies lines from a local terminal to another user. This rather old command was an early form of electronic mail.

The following sections explain how to use these commands, or they point you to the appropriate documentation.

Installing, Configuring, and Using uucp

The `/usr/bin/uucp` facility lets you transfer files among HP-UX and UNIX ™ systems.

Prerequisites and Conditions

- You need two or more systems running HP-UX or UNIX and linked by RS-232-C modem or direct connections.
- Your system is the local system. Any other system is a remote system.
- Meet with system administrators of remote systems to agree on hostnames, the direction of file transfers, and so on.

Procedure

See the *UUCP HP-UX Concepts and Tutorials* manual to get complete information about planning, installing, configuring, customizing, and using `uucp`. The manual also explains related commands (for example, `cu`, `uucico`, `uux`). Besides working manually, you can use SAM as follows:

`System Administration Manager`

↓

`Networks/Communications`

↓

`UUCP Configuration ->`

The following page shows the specific tasks you can perform under this item.

Under `UUCP Configuration ->`, you can select and perform the following tasks:

- `Add a device`
- `Remove a device`
- `View/Modify Device Configurations`
- `Add a System`
- `Remove a System`
- `View/Modify System Configurations`
- `Modify Permissions When Remote Systems Call In`
- `Modify Permissions When Local System Calls Out`

Setting Up mailx and mail or elm

The `sendmail` facility in the NS-ARPA Services provides the `mailx` and `mail` commands. As an alternative, you can set up the `elm` facility.

Prerequisites and Conditions

- HP-UX systems that have been networked.
- Update HP-UX, if necessary, to include the filesets in the NS-ARPA Services.
- For users to get a message saying they have mail, the `.profile` file in the home directory should contain these lines:

```
if [ -f /bin/mail ]
then
        if mail -e
        then    echo "You have mail."
        fi
fi
```

- The code for `csh.login` file is similar. You can get the code from the `/etc/profile` and `/etc/csh.login` files.

Procedure

You can use SAM as described earlier in the module on networking. Otherwise, proceed as follows:

1. To install mail, see the part named "Internetwork Mail-Routing (Sendmail)" in the *Installing and Maintaining NS-ARPA Services.* This document has complete information.
2. To use electronic mail, see:
   a. The chapter named "Sending and Receiving Mail" in the *Beginner's Guide to Using HP-UX.*
   b. The article named "Mailx" in the *Shells and Miscellaneous Tools HP-UX Concepts and Tutorials.*
3. To set up and use `elm`, see *A Beginner's Guide to Using HP-UX.*

Using news

The `news` command lets you place announcements to users on a system. Users get a message indicating they have news during the login process.

Prerequisites and Conditions

- The `/etc/news` directory must contain one or more news files.
- The `.profile` file (or `csh.login`) in the home directory should contain the following code:

  ```
  if [ -f /usr/bin/news ]
  then news -n
  fi
  ```

  You can get the code from `/etc/profile` and `/etc/csh.login`.

Procedure

1. Ensure that `.profile` or `.csh.login` contain the code shown above.
2. Create news files in `/usr/news`, for example `acme-editor`:

   ```
   The ACME Editor was added to the system last night.
   To use the editor, get into a shell and type:

      acme-ed [Return]

   Press ? to get help.  The help explains what to do.
   ```

3. During login, users see a message indicating they have news. To see the news, users execute:

   ```
   news
   ```

4. Each news file appears, preceded by a header such as:

   ```
   acme-editor (root) Mon Apr 3 23:53:04 1989
   ```

5. If a news file scrolls off the display, execute:

   ```
   news | more
   ```

Using wall or cwall

The command immediately broadcasts a message to every user.

Prerequisites and Conditions

- On an HP-UX cluster, `wall` writes to users on the system from which the command executed. On an HP-UX cluster, `cwall` writes a message to users in the cluster.
- The message goes to logged-in users and is preceded by:

  ```
  Broadcast Message from ...
  ```

- In a windowed environment, the message is disruptive and potentially confusing because it appears in every text window including the window from which the user started the window system.

Procedure (Use the same one for cwall in HP-UX clusters)

1. Execute: `/etc/wall`.
2. The cursor drops down one line, waiting for you to type the message such as:

   ```
   The system must go down in one minute.

   Save your work and log off.  Your sys admin apologizes
   for the inconvenience.
   ```

3. When you finish typing the message and wait the appropriate amount of time (indicated in the message), type:

   Ctrl-D

Using write

A person in a group can use `write` to communicate interactively with another user.

Prerequisites and Conditions

- Have a terminal connected to a multi-user system.
- Be a member of a group.
- Terminal users must not have used `mesg` in their `.profile` files to deny write permission.

Procedure

Use the following syntax:

`write` *user* [*message*]

The `write` command sends *message* (if supplied) to the terminal of `user`. Interaction between the users continues until one of the following things occurs:

1. an end-of-file is read from a terminal (^D);
2. an interrupt is sent (esc); or
3. the recipient executes `mesg n`.

See *write*(1) in the *HP-UX Reference* manual to get more information.

Providing for System Security

To administer your system, you may need to deal with security features and policy. Secure systems have scope and depth.

Scope is the ways in which connected and networked systems are secured.

Depth is the ways in which a particular system is secured.

For example, the scope could relate to securing a cluster server in an HP-UX cluster, all the clients in the cluster, and every system that can access the cluster server. Still in this context, the depth could relate to securing the cluster server.

If you need to set up a secured system, you have two areas of concern:

1. Normal system administration security functions; and
2. Security-specific functions that are described in the *HP-UX System Security* manual.

This section describes things you can do as a normal system administrator. For other aspects of having a secure system, see the *HP-UX System Security* manual.

Continue in this section before going to the system security manual.

Conceptualizing Protection and Security

The system administrator (or an official entity) establishes guidelines, goals, strategies, and procedures for preventing illicit activities and inadvertent actions on a system. To accommodate these things, you consider physical, password, permission, and execute security.

Physical Security

- Maintain backup tapes (or other appropriate media) for the system and all layers on the system (environments, applications, users, and such) that could not be replaced from other sources.
- Keep backup media protected and locked in a safe location. This can vary from locking tapes in your desk to placing them in a vault in a remote location.
- Clearly identify LAN and other cables to system users. If appropriate, meet with users and discuss the system, guidelines for dealing with problems, and such.
- Do not leave the root console or any console on which you logged in as the root user unattended. Log off the system anytime you are not physically monitoring it. If necessary, use `lock` to secure the system while you are away.
- If necessary, devise ways to secure hardware against theft.

Password Security

The `/etc/passwd` file lets you secure a system. Although users can read the file, only the root user can alter it. HP-UX encrypts passwords entered by users. You should encourage all users to have a password, and when you add users to the system, force them to provide a password before they can log in. See `/etc/passwd` in "Editing Environment Files" earlier in this chapter if you want details about the password file.

Permission and Execute Security on Files

HP-UX files have modes that determine file access permissions for three categories of users. Assuming you have the string, `-xwrxwrxwr`, the scheme for setting permissions works as follows:

- (dash)	means the file is just a file. Instead of seeing `-`, you see `d` for directories and `c` for special files.
First **xwr**	provides permissions for the owner of the file (a file, directory, or special file).
Second **xwr**	provides permissions for the group, where a group includes the users who need to access the same files.
Third **xwr**	provides permissions for the public (other), where public includes all people who have a system account.
r	means a user can read the file (a dash denies).
w	means a user can write to the file (a dash denies).
x	means a user can execute the file (a dash denies).

The discussion of permissions continues on the next page.

The position of an **x**, **w**, or **r** indicates who can read, write, or execute. Substituting dashes by position lets you control permissions. The following examples suggest possibilities:

`drwxr--r--`	*the owner has all permissions, group and public can read*
`crw-rw-rw-`	*all users can read or write to the special file*
`drwxrwxrw-`	*the owner and group can read, write, and execute files in the directory, the public can read or write to the files*
`-rw-r--r--`	*the owner can read and write, the group and public can read the file*

Use the `/etc/chmod` command to set permissions.

`chmod` *mode filename*

where:

mode is a three-digit octal code for the permissions (for example, `754`); and

filename is the file or directory getting the permissions.

A Typical Set of Protections

- Set the root directory to 555, (`dr-xr-xr-x`).
- Set `/etc/passwd` to 444, (`-r--r--r--`).
- Set the default file creation mode so new files have write protection to all but the owner. Use `umask` with a mode of 022 (`-----w--w-`) and include the entry in `/etc/profile` and `/etc/csh.login`.
- Set system directories `/usr`, `/lib`, `/usr/lib`, `/bin`, `/usr/bin`, and `/etc` to 555 (`dr-xr-xr-x`) to prevent users from adding or deleting files.
- Set temporary directories such as `/tmp` and `/usr/tmp` to 766 (`drwxr-xr-x`).

Security for the Root User

You can secure login access for the root user by:

- Setting a password for root in the usual way (login as root and execute `passwd`).
- Specifying which terminal can be used to log in as the root user.

For the second method, work as follows:

1. Specify the tty devices the root user can use for logging in by creating `/etc/securetty`. Unless you explicitly create this file, the root user can log in from any terminal connected to the system. Create the file as follows:
   a. Entries in the file specify the device file names of the terminals where the root user can log in.
   b. The entries contain the name of the device file for the terminals, but not the pathname (typically `/dev`).
   c. The security file can specify more than one terminal, one to a line.
   d. In an HP-UX cluster, the file must be a Context Dependent File (CDF).
2. A security file that permitted the root user to log in on the terminals having the device files for `console` and `tty05` would look like this:

```
console
tty05
```

Limitations in Creating /etc/securetty

This security feature does not restrict a normal user from using `su` to become the root user on any terminal. You can restrict the use of `su` by executing:

```
chmod 500 /bin/su
```

Then, the owner (root) has read and execute permission, and other users have no access permission.

HP-UX Cluster Security

In general, you secure an HP-UX cluster just as you secure a standalone workstation. But the situation is magnified.

To prevent problems, require every cluster client and remote user to log in and provide a password.

Some General Tips for Securing a System

Earlier, you saw information for physical, password, and root security. This section contains general tips learned by expert system administrators that help you secure a system.

The following items mention basic principles of good security:

- Physically control equipment.
- Get management commitment to security.
- Educate employees on what is expected of them.
- Use administrative procedures designed to increase security.
- Concealment alone is not security.
- Do not trust what others can alter.
- It is better to know about holes in security than to think you are secure.
- Compartmentalize data.
- Disconnect unused terminals and mass storage devices.

Routine Tasks That Help You Maintain Security

- Check `/usr/adm/sulog` and `/usr/adm/OLDsulog` periodically for successful and unsuccessful attempt to *su*(1).
- Check `/etc/btmp` periodically for unsuccessful logins.
- Check log files and directories in `/usr/adm`. Check log files in the directories under `/usr/spool`.
- Check the log files for any specific software you have on the system.
- Do not put a period (.) in the path of the root user.
- Check out each setuid program as to why they are setuid.
- All accounts should have a password, or they should have * as the password.
- A modem line has the potential to be a security hole.
- Protect `/usr/lib/uucp/Systems`, `$HOME/.netrc`, any scripts that use RFA and NFT commands, and any other files that have passwords of remote systems embedded in them.
- Use password aging.
- Use C2 Trusted Systems and ACLs.
- The following items show files to protect (owners and permissions):

```
/dev/kmem   root 400
/dev/mem    root 400
/dev/rhd (300)   root 400
/dev/dsk/*   root 400
superuser's login directory root 555
```

Procedures for Setting Up a Secured or Audited System

The *HP-UX System Security* manual has information about procedures. Also, you can use SAM to perform several tasks related to auditing and security by selecting the following menu options:

`System Administration Manager`

↓

`Auditing and Security - (Trusted System) ->`

↓

You can then perform the following tasks:

- `Turn Auditing On/Off`
- `Set Audit Monitor and Log Parameters`
- `Convert to Trusted System`
- `View Audit Logs`
- `View/Modify /what is Being Audited: Users`
- `View/Modify /what is Being Audited: Events`
- `View/Modify /what is Being Audited: System Calls`

System Accounting

You may need to do some system accounting. The amount and type of accounting depend on your type of system and the external policies that affect your administration.

Whatever your situation may be, if you need to do system accounting, see the appropriate chapter in the *HP-UX System Administration Concepts Manual.* That chapter explains concepts, guidelines, and procedures. In addition, it provides some reference information.

Except to discuss some log files, the information about accounting does not appear in this manual because you may not need to do much accounting.

3

Updating HP-UX

This chapter describes how to update Series 300 and Series 800 HP-UX systems to HP-UX release 7.0. The table suggests how to use this chapter and recommends alternate references.

To Do This ...	**Work As Follows ...**
Update to HP-UX 7.0 from 6.2/2.1 or earlier.	■ Update to 6.5/3.0; then use this chapter and the online help screens. ■ Or, back up your system, and then re-install using the HP-UX installation manual for your system.
Interactively update the HP-UX operating system and core product files to 7.0.	Use this chapter and the online help screens. (You must have HP-UX release 6.5, 3.1, or 3.0.)
Interactively update HP-UX with new or existing optional applications.	Use this chapter, the online help screens, and the documentation accompanying the application software.
Perform a non-interactive update.	Refer to *update*(1M) in the *HP-UX Reference* and "Running Update Non-interactively" in this chapter.
Convert a system into a network software distribution server.	See "Configuring a Netdist Server" in this chapter or *netdistd*(1M) in the *HP-UX Reference.*
Convert an existing Series 800 to be a cluster server.	Re-install. See "Managing an HP-UX Cluster" in this manual.
Move to a more secure operating system.	Re-install. See "Setting Up and Maintaining Your Secure System" in the *HP-UX System Security Manual.*

New Features in 7.0 Update

This module briefly describes how the 7.0 `update` program differs from previous HP-UX update programs.

The 7.0 `update` program:

- Has a new interactive user interface with these features:
  - Lets you select items from menus.
  - Provides online help.
  - Uses function keys extensively.
  - Accepts no arguments at the command line (you can use arguments *only* when updating non-interactively).
- Runs on both the Series 300 and Series 800. The update process is essentially identical; however, you might notice these differences:
  - The default source for the Series 300 is `/dev/update.src`; for the Series 800, the default is `/dev/rmt/0m` (i.e., *zero*m).
  - The Series 300 typically uses cartridge tapes; the Series 800 uses 9-track tapes or cartridge tapes.
  - The Series 300 runs unattended once you start loading (except for changing tapes); the Series 800 might require that you answer some questions during the load. Then the process runs unattended.
- Calculates disk space requirements before it begins loading the software. This feature prevents loading unless you have sufficient space on your destination disk. You can adjust your selections and disk space before loading begins.

- No longer reboots immediately (*Series 800*) or requires single-user state. Many filesets (primarily those for applications) can be updated without disturbing other users. Use single-user state if your update will affect the operating system (i.e., the kernel).
- Lets you update your system from a **netdist server** (a system on the network configured for network distribution). To configure a netdist server, see "Configuring a Netdist Server" in this chapter or *netdistd*(1M) in the *HP-UX Reference*.
- Lets you select the new filesets that match those already on your system. (Fileset and partition names and contents have changed for 7.0; if you select this option, all files might not be updated.)

An Overview of the Update Program

You can run the `update` program in two ways:

- Interactively (using menu-driven screens and online help).
- Non-interactively (using command line options).

Running Update Interactively

Throughout the interactive update process, menu items or function keys let you perform the following tasks (online help is available to explain these tasks):

- Update a local system from tape or from a netdist server. (When updating from a netdist server, you must have already configured the server for network distribution. For more information, see "Configuring a Netdist Server" in this chapter or *netdistd*(1M) in the *HP-UX Reference*.)
- Change the source or destination, if necessary.
- Load all partitions and filesets from the source media.
- Load only the filesets that exist on both the source media and destination disk. (Fileset and partition names and contents have changed for 7.0; if you select this option, all files might not be updated.)
- Select or de-select individual partitions and/or filesets to load from the source media to the destination disk.
- Analyze disk space requirements, if desired. The `update` program automatically calculates disk space requirements before it begins loading filesets. HP-UX 7.0 is larger than 6.5/3.1; therefore, when you update HP-UX or applications, a portion of your free disk space might be consumed. You can analyze disk space before loading to see how your selections will affect this free space. (See "Analyzing Disk Space" in this chapter. For a detailed discussion of disk space, refer to the *HP-UX System Administration Concepts Manual.*)
- Exit the `update` program, or return to the main menu at any time (except once you start loading).

If Problems Occur . . .

Many errors can be handled immediately during an interactive update. If an error occurs during the `update` program, a message prints telling you what happened. This message (together with any other error messages) is recorded in `/tmp/update.log`. For additional information, see "Checking the update.log File" and "A Guide to Troubleshooting" in this chapter.

Running Update Non-interactively

The non-interactive (command line) interface is less friendly than the interactive `update` program. You might want to use the non-interactive interface under these conditions:

- You have a non-HP terminal that is incompatible with the screen control used in the interactive version of `update`.
- You use `update` regularly and want to bypass the interaction.

When updating from the command line, you provide at invocation all the information you would be prompted for interactively. Refer to *update*(1M) in the *HP-UX Reference* for descriptions of command line options.

If Problems Occur . . .

Errors that occur *before* loading begins produce messages to the file `/tmp/update.log` and to standard error. These errors cause `update` to quit immediately.

Errors that occur *after* loading begins are logged to `/tmp/update.log` and written to standard error. Review these errors after loading completes.

For additional information, see *update*(1M) in the *HP-UX Reference*. Also see "Checking the update.log File" and "A Guide to Troubleshooting" in this chapter.

Analyzing Disk Space

The `update` program determines if you have sufficient disk space before it begins loading filesets. The amount of disk space available on your destination depends on these factors:

- The files on your system.
- The size of your disk (or disks).
- The partitions and filesets you select to load.
- The minimum free space required on your system's mounted file systems.
- Swap space requirements on your system.
- Future expectations for space requirements on your system. To deal with space problems over the long term, you might consider mounting an additional disk device or adjusting swap space.

The table below gives you some guidelines for the amount of disk space needed if you're updating from 6.5/3.1 to 7.0 HP-UX or if you're updating a mixed cluster.

If You Have ...	**You Need ...**
Series 300	■ 14 additional Mbytes for AXE, which includes the X11 and Starbase run environments. ■ 12 additional Mbytes for PE, which includes the remaining X11 filesets; Starbase, FORTRAN, and Pascal libraries; and support for heterogeneous clusters. ■ Also consider space needed for additional applications.
Series 800	■ 22 additional Mbytes for hpuxcore. ■ Also consider space needed for additional applications.
Mixed cluster	■ Enough space for Series 300 filesets *and* Series 800 filesets. For 7.0, AXE and PE require 83 Mbytes, total; hpuxcore requires 75 Mbytes, total. ■ Also consider space needed for additional applications. A mixed cluster server, for example, probably will need at least 400 Mbytes of space—plus space for future growth.

Dealing with Minimum Free Space Requirements

You must have a certain amount (10% by default) of free disk space on each mounted file system to which you will load files. This space is called "minimum free space," or *minfree*. Only the superuser (a person logged in as `root`) can allocate space on a file system with less than *minfree*.

As you run `update`, you might see one of these messages:

```
It is recommended you free up n kbytes

Loading the selected filesets will result in less free disk space ...
```

Either message means that less than *minfree* will remain on one or more file systems to which you will load files. If you are using `update` interactively, you can proceed with loading. However, it is recommended that you first free the minimum free space (see "How to Free Disk Space" on the following page).

Dealing with Insufficient Disk Space

As you run `update`, you might see one of these messages:

```
You MUST free up n kbytes

Loading the selected filesets is impossible due to insufficient space on one
or more file systems....
```

Either message means that you do not have sufficient disk space on one or more file systems to complete your update, and you cannot proceed with loading. You must free a sufficient amount of the file system's disk space before you can proceed with the program (see "How to Free Disk Space" on the following page).

How to Free Disk Space

You can free disk space by de-selecting filesets or removing files, by mounting another file system, or by creating symbolic links.

De-select Filesets or Remove Files Interactively

While in the `update` program, you can free disk space by de-selecting filesets using the Fileset Selection screen, the Disk Space Analysis screen, or the Partition Selection screen.

Alternatively, you can remove unnecessary files from the file system by pressing the `Shell` function key, removing files, and returning to the `update` program. To inspect files larger than *n* kbytes, use the **`find`** command with a size parameter (see *find*(1) in the *HP-UX Reference*).

The specific files you might want to remove will vary. To start, however, check the following for unnecessary files: `/tmp`, `/etc/*tmp*` (you might remove `wtmp` or `btmp`), `/usr/adm`, `/usr/local`, `/usr/contrib`, `/usr/tmp`, `/users`. You can remove unnecessary filesets using the `sysrm` command (see *sysrm*(1M) in the *HP-UX Reference*).

Note Do not remove files or directories under `/system`. If you do, you won't be able to convert to (or create) an HP-UX cluster. Also, system customization may fail if you decide to load new filesets later.

Mount Another File System

The chapter in this manual titled "Managing the File System" describes how to mount additional file systems to free disk space.

Create Symbolic Links

You can free disk space by moving files or directories and creating "pointers" (or **symbolic links**) from the old locations to the new locations. Symbolic links can span file systems and refer to directories as well as files. Create symbolic links with the command `ln -s` (see *cp*(1) in the *HP-UX Reference*).

For example, to move `/usr/man` from `/usr` to `/extra/man` (where `/extra/man` is a file system in which more free space exists), follow the steps below. Because `/usr/man` is approximately 5 Mbytes, you will be freeing 5 Mbytes in the `/usr` file system and consuming that amount in `/extra`:

1.	`cd /usr` `find man -print | cpio -pdumv /extra`	*1. Copy the subdirectory from /usr to /extra.*
2.	`mv /usr/man /usr/man.old`	*2. Temporarily rename the original man (this is your "backup").*
3.	`ln -s /extra/man /usr/man`	*3. Create a symbolic link between the directory's new location and its old location.*
4.	`ls /usr/man`	*4. List the contents of the directory.*
5.	`ls /usr/man.old`	*5. List the contents of your "backup"; the output should match the output of Step 4.*
6.	`rm -rf /usr/man.old`	*6. If the outputs match, remove your "backup."*

If the outputs from Step 6 *don't* match, remove the link, and start over at Step 1. To remove the link, type the following:

```
rm -rf /usr/man /extra/man
mv /usr/man.old /usr/man
```

Getting More Information

- Refer to *update*(1M) in the *HP-UX Reference* to learn more about using `update` non-interactively.
- Refer to *sysrm*(1M) in the *HP-UX Reference* to learn more about using the `sysrm` command.
- Refer to *cp*(1) in the *HP-UX Reference* to learn more about using the `ln` command.
- Refer to the *HP-UX System Administration Concepts Manual* for a more detailed discussion of swap space, disk space requirements, and *minfree*.

Planning for an Update

This module describes what you must do or know *before* you begin updating your system.

Prerequisites and Conditions for All Updates

- Make sure you have adequate disk space to contain 7.0 HP-UX. For information on 7.0 disk space requirements, and on analyzing and freeing disk space, see "Analyzing Disk Space" in this chapter. For a detailed discussion of disk space, see the *HP-UX System Administration Concepts Manual*. For instructions on configuring mixed clusters, see "Managing an HP-UX Cluster" in this manual.
- *Series 800.* If you intend to use disk mirroring, make sure your system can accommodate it. For additional information, see the *Series 800 HP-UX System Administration Tasks Manual.*
- You must know the correct device file name for your cartridge tape drive or your 9-track tape drive. On a Series 300, the `update` program defaults to `/dev/update.src`; on a Series 800, the program defaults to `/dev/rmt/0m` (i.e., *zero*`m`). However, the device file names can vary: if your system does not have the default device file, or if you will be updating from a different device or from cartridge tape (on Series 800), determine the name of the source device's device file name now.

  If you don't know the device file name, determine its name from the information given in a directory listing:

  - *Series 300:* Type one of the following:
    - `ll /dev` (if all your device files are in the `/dev` directory).
    - `ll /dev/rct` (if your device files are grouped in sub-directories).

    For information on interpreting the device file listing, see "Managing Devices" in this manual.
  - *Series 800:* If you are updating from 9-track tape, type `lssf /dev/rmt/*` to view your device files. If you are updating from cartridge tape, follow this procedure:

    1. Determine your tape drive's hardware address.

2. Change directory to `/dev/rct`.
3. Type: `lssf *` (see *lssf*(1M) in the *HP-UX Reference* for more information).
4. A listing of all your device files is displayed. Choose the device file that matches your hardware address.

- Read the *Read Me First* document supplied with your update media.
- Back up your system before you begin a system update (for information on this procedure, see the chapter in this manual titled "Backing Up and Restoring the System").
- Save or back up your kernel in `/SYSBCKUP`.
- Run `/etc/fsck` on all disks (*Series 300*) and/or all disk partitions (*Series 800*). For additional information on this command, see *fsck*(1M) in the *HP-UX Reference* and "Managing the File System" in this chapter.
- Set up a netdist server if you plan to update over a network (see "Configuring a Netdist Server" in this chapter).

  If a netdist server is already set up, know the hostname of the netdist server from which you will update, and know the port number. Find the port number in the `/etc/services` file on the netdist server (typically, this number is 2106).
- If you are on a cluster, your media must be on a drive connected to the root server, and you must execute the `update` program from the root server.
  - If you are on a Series 800 cluster, first complete the Series 800 portion of the update; then complete the Series 300 portion. A Series 800 cluster server must use Series 300 media to update the files for Series 300 cnodes.
- If you are using an autochanging device as your source drive, use the switch on the back panel to set the unit to "sequential" mode. Load the update tapes in sequential order in the autochanger magazine starting with slot 1 (bottom slot). Then load the magazine in the drive and wait until the drive is ready.

More "Prerequisites and Conditions" on the following page....

- The `/etc/update` facility will mount all file systems listed in `/etc/checklist`. If you don't want a file system to be mounted (for example, `/users`), comment out the line in `/etc/checklist` before you start the `update` program. Leave NFS file systems (if any) in the checklist file.
- If an NFS server accesses any portion of your operating system, update the server system first.
- Be aware that fileset (and partition) definitions and names have changed for the 7.0 release; thus, fileset contents may differ between 7.0 and previous releases. It is possible that not all obsolete files will be removed. Also, if you choose to update only the filesets that exist on both the source and destination, it is possible that all files might not be updated. (The *HP-UX System Administration Tasks* manual have information about partitions and filesets.)
- Be aware that the `update` program might reboot the system as part of the update process. The program reboots the system if you select a fileset (or filesets) that are flagged to indicate a reboot is necessary. If a reboot is necessary, you will be warned before loading begins. You then have the option of proceeding or exiting the `update` program, bringing the system to a quiet state, and re-starting `update`.

  For additional information on system reboots, see *update*(1M) in the *HP-UX Reference*.
- Ensure that your TERM environment variable is set correctly; otherwise the display might behave strangely during the interactive update process (setting the TERM variable is described in *A Beginner's Guide to Using Shells*). You can use CTRL-L to refresh the screen provided `update` isn't loading filesets.
- Some HP terminals have a timeout that blanks the screen after a period of keyboard inactivity. To refresh the screen after a timeout, press the Shift key; this key is not interpreted as input.
- If you are using a non-HP terminal, you cannot run `update` interactively. Refer to *update*(1M) in the *HP-UX Reference* for instructions on running `update` using command line options.
- Don't run console jobs in the background during an update: unexpected output to the terminal might result in unreadable update screens.

Prerequisites and Conditions for Updating the Operating System

- Read and follow the prerequisites and conditions on the previous pages.
- Make sure users are logged off the system.
- The `/hp-ux` file must have been built from `/etc/conf/dfile` (*Series 300*) or `/etc/conf/gen/S800` (*Series 800*). Also, the system must have been most recently booted from `/hp-ux` (*not* from `/SYSBCKUP`). If this is not the case, follow this procedure:
  1. Move `/etc/conf/dfile` to `/etc/conf/dfile.old` (*Series 300*), or move `/etc/conf/gen/S800` to `/etc/conf/gen/S800.old` (*Series 800*).
  2. Move the file that was used to create the `/hp-ux` file into `/etc/conf/dfile` (*Series 300*) or into `/etc/conf/gen/S800` (*Series 800*).
- If the system is not currently running on `/hp-ux`, you might want to reconfigure the kernel and reboot the system on `/hp-ux` (see "Reconfiguring the Kernel" in this manual).
- Shut down the cnodes if updating an HP-UX cluster.
- Bring the system to single-user state (see *shutdown*(1M) in the *HP-UX Reference*). If you have remote file systems mounted via NFS that will be updated, these file systems must be mounted during the update. NFS must be active to detect which files are remote; however, `update` will not write to NFS-mounted file systems. (To use NFS or to update from a netdist server, bring up the system in multi-user state, then shut down to single-user state.)

Where to Go Next ...

If You Plan to ...	Move on to Section ...
Update from cartridge tape, 9-track tape, or netdist server.	Work through the module "Running the Update Program." Then move to "Tasks to Perform after the Update."
Set up a netdist server.	Work through the module "Configuring a Netdist Server." Then, if you want to perform an update, move to "Running the Update Program."

Running the Update Program

Prerequisites and Conditions

- Read the module "Planning for an Update" in this chapter.

Procedure

1. Become the root user on the system you want to update.
2. If you are updating the operating system, or loading any filesets that will cause the system to reboot, bring your system to a quiet state now.
3. You have three cases for source media. Find your case in the table on the facing page, and take the appropriate action. Then go to Step 4.

Note You must use the `tar` command to get the 7.0 update program from tape, as indicated in the table. Know the device file. If it is different from the one indicated in the table, substitute your device file. For more information on device files, see "Planning for an Update" in this chapter.

4. After executing the `tar` command, wait several minutes for the first `TOOLS` file to be extracted; once the process begins, these files are extracted rapidly. When the last `TOOLS` file is echoed to the display, wait 5-7 minutes, and press (BREAK) to end the tar process (`TOOLS` files are listed in Appendix C).
5. Make sure you're in the root directory (`cd /`). (If updating from a netdist server, remove `/tmp/TOOL`.) Then, type `/etc/update` to invoke the interactive `update` program.
6. The Main Menu appears (see the next module). From this point, the update process is interactive: select items from a menu, and use the online help screens to explain specific options. After making your selections, begin loading (do NOT touch the keyboard during the load process):
   - *Series 300:* The load completes unattended (except for changing tapes).
   - *Series 800:* You may need to answer questions during the load.

If Updating from ...	Do This ...
Cartridge tapes	The protect switch must point to `SAFE` or `PROTECT`. Insert the first tape; proceed when the BUSY light remains off. At the system prompt (`#`), type one of the following: ■ *Series 300:* `tcio -iZ /dev/update.src | tar -xvf - TOOL` ■ *Series 800:* `tcio -iZ /dev/rct/c3d0s2 | tar -xvf - TOOL` (NOTE: Device files may vary. See "Planning for an Update.")
9-track tapes	Ensure that the write ring has been removed, mount the tape, and put the tape drive online (the ONLINE light is on). At the system prompt, type: `tar -xvf /dev/rmt/0m TOOL` (i.e., *zero*m) (NOTE: Device files may vary. See "Planning for an Update.")
Netdist server	Have a server configured for network distribution, and have the `update` program available to clients (see "Configuring a Netdist Server" in this chapter.) Copy the 7.0 TOOL fileset from the configured netdist server to the system you want to update (for information on using/setting up anonymous ftp, see *sam*(1M) or *Installing and Maintaining NS-ARPA Services*): 1. Type: `ftp` *netdist_server* (*netdist_server* is the name of the server from which you are copying `update`). 2. At the login prompt, type: `ftp` 3. At the password prompt, type: `anonymous` 4. You should see the prompt `ftp>`: ■ *Series 300.* Type: `get dist/TOOL.300 /tmp/TOOL` ■ *Series 800.* Type: `get dist/TOOL.800 /tmp/TOOL` 5. You should see messages similar to these: `Opening data connection for dist/....` *n* `bytes received....` 6. When you see the `ftp>` prompt again, type: `quit` 7. At the local system prompt, type: `cd /` 8. Then type: `tar -xvf /tmp/TOOL`

Using the Main Menu

When you invoke the interactive `update` program, you see the Main Menu:

```
UPDATE                              Main Menu

      Highlight an item and then press "Return" or "Select Item".
      To refresh the screen press CTRL-L.

      From Tape Device to Local System

      Source: /dev/update.src (Series 300)               Destination: /
              /dev/rmt/0m     (Series 800)

               Load Everything from Source Media ->
               Update Only the Filesets on the Destination ->
               View or Select Individual Partitions  ->

               Change Source or Destination ->

               How to Use Update

--------------------------------------------------------------------------------
 Help  |        | Shell  | Select |        |        |        |        |  Exit
       |        |        |  Item  |        |        |        |        | Update
--------------------------------------------------------------------------------
```

Navigate in `update` the same way you navigate in SAM. For details on navigating and using soft keys, see "Using the System Administration Manager" in chapter 1 of this manual, or you can highlight the "How to Use Update" menu item, and press `Help` (f1).

During the update process, you will progress through a series of screens. If you are uncertain about the function of certain menu items or data entry fields, highlight that item or field, and press `Help`. When you've completed all selections, initiate the loading process. The `update` program always prompts you for confirmation before it begins loading.

Selecting "Change Source or Destination" from the Main Menu

Select "Change Source or Destination" under these conditions:

- Your destination or source differs from the default listed on the Main Menu. The default destination for `update` is `/` (`root`); change destination only for certain optional applications. The default source depends on system type and update media:
  - *Series 300:* Default source device file is `/dev/update.src`.
  - *Series 800:* Default source device file is `/dev/rmt/0m` (i.e., *zero*`m`).

  However, if you use cartridge tape instead of 9-track tape, your source will likely be `/dev/rct/c3d0s2`; change the default source to agree with the name of your source device's device name (see "Planning for an Update").
- You are updating via netdist server.

This table explains how to change the default source or destination:

Updating from Local Source	Updating from Netdist Server
1. Select "Change Source or Destination" from the Main Menu.	
2. A pop-up menu appears. Select "From Tape Device to Local System".	2. A pop-up menu appears. Select "From Netdist Server to Local System".
3. A form appears. If you have not already created a device file for the device from which you will update, enter `y`, and specify the new address for `/dev/update.src`. Otherwise, enter `n` and then the name of the source device file from which you will update. You also can change the destination from this form (change the default destination only for certain optional applications).	3. A form appears. Enter the netdist server name (either the netdist server's host name or internet address; if you don't know this name, see your system or network administrator). You also can change the port number and/or destination from this form (the default port number, 2106, typically is correct).
4. Press `Done` (f4). The Main Menu re-appears. Now choose from the list of menu items (to learn more about a menu item, highlight it, and press `Help`).	

Tasks to Perform after the Update

After the `update` program has completed, you need to perform various tasks to set up the system for users. These tasks include the following:

1. Log in and check the following files and directories:

Check this ...	For this Information ...
`/tmp/update.log` file	Contains a description of the events and any errors that occurred during the update process. For more information, see "Checking the update.log File" in this chapter.
`/etc/newconfig/Update_info` directory	Contains files with software product update information.
`/etc/newconfig/README` file	Contains useful information about files in `/etc/newconfig`.
`/etc/newconfig` directory	Contains new versions of some files normally put into the `/etc` directory (for example, `rc`, `brc`, `backup`, `backupf`). Because you might have edited the original versions of these files, they are not replaced by either `update` or customize scripts. You can use the `diff` command to find the differences between old files and the new files placed in `/etc/newconfig`. Incorporate your changes into the new files, move them to `/etc`, and use the new files.

2. *Series 300:* Create a new recovery system (see "Constructing an HP-UX System" in this manual).
3. If you have a standalone system, you might want use SAM to set up a cluster environment and add cnodes to your cluster (for how to perform these tasks, see "Managing an HP-UX Cluster" in this chapter).
4. If you want to convert your system to a cluster server, see "Managing an HP-UX Cluster" in this chapter.
5. If you have a Series 800 cluster server, and you want to service Series 300 cnodes, follow these steps:

   a. Update the Series 800 using Series 300 media.
   b. Use SAM to add Series 300 cnodes.

Checking the update.log File

After running the `update` or `updist` programs, check `/tmp/update.log` for a description of the events and any errors that occurred during the process. This log file has a time and date heading that corresponds to each update session. After you've reviewed this log, you might want to remove it: the `update` program will continue to append to this log file, and it could become large.

The `update.log` file contains three broad kinds of messages:

1. Output from `update` or `updist`. These messages are indented nine spaces and might be preceded by a message label:

`=======`	Indicates that a task is beginning or has completed. For example:
	`======= ... BEGINNING UPDATE PROGRAM`
	`======= ... COMPLETED UPDATE PROGRAM`

`ERROR:`	Indicates the program cannot proceed, or that it needs corrective action. For example: `ERROR: Destination directory "/mount" is invalid: No such file or directory.`
`WARNING:`	Indicates the program can continue. However, something went wrong or requires attention, either now or later. For example: `WARNING: Cannot access /etc/checklist file: No such file or directory.`
`NOTE:`	Indicates that something out of the ordinary or worth special attention has happened. For example: `NOTE: Saved /hp-ux in /SYSBCKUP.`

(No prefix)	Indicates generic progress and status messages. For example:
	`Beginning to load fileset "KERN_BLD".`
	`Successfully loaded fileset "KERN_BLD".`

All `ERROR` and `WARNING` messages are also displayed to the screen, whether `update` or `updist` runs interactively or from the command line.

2. Output from fileset customize scripts, which should be in the same format described above.

3. Output (standard output or standard error) from programs executed by `update`, `updist`, or a customize script. For example, the following output is from a program executed by `update`:

```
         Following is output from "/bin/make":
================================================================================
Compiling conf.c ...
        /bin/cc  +O1 +M -Wc,-Nd3500,-Ns3500 -Dhp9000s200 -D_hp9000s300 -D_hp90

 ...
```

A Guide to Troubleshooting

The table in this module describes some of the possible situations and error messages that you might encounter during the interactive update process.

Situation	Possible Message(s)	What to Do
Updating from a local source.	`... The specified source could not be verified. Change the source specified, or ensure the media is loaded and ready to read.`	1. Ensure either the 9-track drive is online or the cartridge drive is finished loading the tape (the busy light is out). 2. Ensure your source is correct: ■ *Series 300:* Default source is `/dev/update.src`. ■ *Series 800:* Default source is `/dev/rmt/0m` (i.e., *zero*m). If your system does not have the default device file, if you are updating from a different source, or if you are updating a Series 800 from cartridge tape, you must specify the correct name of the source device's device file. See "Planning for an Update" in this chapter.

Situation	Possible Message(s)	What to Do
Updating from a local source.	`Cannot open source` *name*:	Ensure your source is correct: ■ *Series 300:* Default source is `/dev/update.src`. ■ *Series 800:* Default source is `/dev/rmt/0m` (i.e., *zero*`m`). If your system does not have the default device file, if you are updating from a different source, or if you are updating a Series 800 from cartridge tape, you must specify the correct name of the source device's device file. See "Planning for an Update" in this chapter.

This table is continued on the following page....

Situation	Possible Message(s)	What to Do
Updating from a netdist server.	1. `Netdist server host` *name* `is not in the hosts database.` 2. `Cannot connect to netdist server running on host` *x* `at port` *y:* `Server not running or connection refused.` 3. `Connection closed by netdist server on host` *name:* `Access to this server is restricted.` 4. `... Server has no filesets for this architecture.`	For each of these messages, first check that the host name and port number are correct. If you do not know the correct name/number, see the network or system administrator. 1. Either add an entry for the server host to the hosts database, or use the server host's internet address instead of its host name. 2. The netdist server probably is not running on the specified host. Contact the network administrator. 3. The network administrator must modify the server's security file so your system can use the netdist service. 4. The server specified is not distributing software for your system type. Either locate a server that can distribute software for your system type, or configure the server specified so it can distribute the appropriate software.

Situation	Possible Message(s)	What to Do
Updating a cluster.	1. `The Series 300 software on this update media cannot be loaded correctly on a Series 800 clustered system due to missing CDF information.` 2. `The software on this update media might not function correctly on a clustered system due to missing CDF information.`	1. On a Series 800 cluster, `update` requires CDF information that does not exist on pre-7.0 media. Obtain a 7.0 version of this software. 2. If you are planning to convert this system into a cluster server, you need a 7.0 version of this software. If you do not intend to convert the system into a cluster server, you can ignore this message.
Configuring a netdist server (running `updist`).	`The software on this update media cannot be distributed by a netdist server due to missing CDF information.`	Network distribution requires CDF information that does not exist on pre-7.0 media. Obtain a 7.0 version of this software.
Starting the `netdistd` program.	1. `netdist/tcp is not a registered service` 2. `Address already in use`	1. Either add an entry to the `/etc/services` file, or specify a port number using the `-P`*port* option (see *netdistd*(1M) in the *HP-UX Reference* and *services*(4) in the *Networking Reference*). 2. Another netdist server program is running on the specified (or default) port. Specify a different port using the `-P`*port* option.

This table is continued on the following page....

Situation	Possible Message(s)	What to Do
You've inserted the first update tape.	1. `... The media type appears to be Logical Interchange Format (LIF)....` 2. `... The media is an unrecognized (invalid) type, or it might be corrupt.` 3. `... The information files on the media did not exist or could not be read....`	1. You inserted a tape not in 7.0 update format (Series 800 install, Series 300 pre-7.0 update). See *update.6.5*(1M) in the *HP-UX Reference*. 2. Perhaps you inserted a tape not in 7.0 format. See *update.6.5*(1M) in the *HP-UX Reference*. If you suspect corruption, call HP. 3. Perhaps you inserted a tape not in 7.0 format. See *update.6.5*(1M) in the *HP-UX Reference*.
Destination disk is almost full.	■ `It is recommended you free up` *n* `kbytes....` ■ `Loading the selected filesets would result in less free disk space ...`	See "Analyzing Disk Space" in this chapter. For instructions on mounting file systems to free disk space, see "Managing the File System" in this manual. For a detailed discussion of disk space, see the *HP-UX System Administration Concepts Manual.*
Not enough disk space to complete the update.	■ `You MUST free up` *n* `kbytes.` ■ `Loading the selected filesets is impossible due to insufficient space on one or more file systems....`	See "Analyzing Disk Space". For mounting file systems to free disk space, see "Managing the File System". For a discussion of disk space, see the *HP-UX System Administration Concepts Manual.*

Situation	Possible Message(s)	What to Do
You selected filesets that caused a system reboot, but unexpected processes are running.	`... non-essential process (other than those expected at the time of system reboot) ...`	You have several options: ■ Continue, and the processes will terminate eventually, using disk space while `update` runs; users may lose work in progress. ■ Do a shell escape, and run `cat` or `more` on `/tmp/update.procs` to get PIDs of nonessential processes. Kill those processes (`kill -9` *pid*). ■ Do a shell escape and run `ps` to inspect all processes running on the system. ■ Exit `update`, and run `shutdown`; then restart the program by typing `/etc/update`). You will lose any partition/fileset selections already made.

Situation	Possible Message(s)	What to Do
Series 800: The `update` program quit without successfully building a kernel.	You might not see any messages on the display. Messages will be logged to `/tmp/update.log`.	Probably a mismatch between your hardware configuration and `/etc/conf/gen/S800`. Fix `S800`. See "Reconfiguring the Kernel" in the *HP-UX System Administration Tasks* manual. Execute: `/etc/uxgen /etc/conf/gen/S800` to ensure that the file has no typing errors. Type `exit` to return to the update facility, and that facility will attempt to rebuild the kernel.
Series 800: You are given a shell and instructed to fix a gen file.	Messages will vary.	Perhaps due to a mismatch between your hardware configuration file and the S800 file (`/etc/conf/gen/S800`). See "Reconfiguring the Kernel" in the *Series 800 System Administration Tasks Manual;* and perhaps the the hardware installation and configuration guide. Before you re-execute `/etc/update`, ensure that the running kernel, `/hpux` file and kernel generation file agree (see prerequisites).
A customize script has failed.	`Type "exit" to return to update.`	Check `/tmp/update.log`, and re-run the customize script as indicated.

Situation	Possible Message(s)	What to Do
Series 800: Messages are overwriting the update screen (update media: 9-track tape).	Any variety of kernel messages.	Use CTRL-L to refresh the screen if `update` is not loading filesets (do not touch keyboard during loading).

Series 800: The process quits at boot-up, either during or after an update.	`Utility requires more memory than is configured.`	Reset the ISL *fastsize* parameter:
		1. Cycle power on the processor (turn the processor off and then on again).
		2. When you see this prompt, interrupt the autoboot process by pressing any key:
		`To override, press any key within 10 seconds.`
		3. Press `y` at this prompt:
		`Boot from primary boot path`
		4. Press `y` at this prompt:
		`Interact with IPL`
		5. When you see the `ISL>` prompt, type:
		`fastsize f`
		6. Cycle power on the processor again. The autoboot should now succeed.

Garbled display or odd behavior after the update.	No error message	You may have set `TERM` incorrectly. Setting `TERM` is described in *A Beginner's Guide to Using Shells.* CTRL-L refreshes the screen.

Configuring a Netdist Server

Unless you are configuring or administering a netdist server, you can skip the rest of this chapter.

A Brief View of the Situation

The term **netdist** refers to a feature of the 7.0 `update` program that lets you update from a specially-configured server ("netdist server") instead of updating from the product tapes. The netdist service is particularly useful for large groups of networked systems that must be updated to a new HP-UX release.

A **netdist server** is a system on your network that has been set up to distribute HP-UX software and applications over the network. When the netdist server is operational, other systems on the network can specify it as the source during update sessions. The `updist` program (a variant of `update`) transfers filesets from tape to the netdist server's disk. The netdist server can then offer these filesets to other systems over the network.

The `netdistd` (daemon)program runs on the netdist server. This program lets other systems run `update` by drawing filesets from where `updist` deposited them and delivering the filesets over the network to the requesting system. Thus, updating from a netdist server is more convenient than updating from tape and requires no tape drive on the requesting system.

Prerequisites and Conditions

- Have the appropriate networking hardware and know how it is configured.
- Choose a server with ample disk space. For HP-UX 7.0, you need a *minimum* of 83 Mbytes for AXE and PE (Series 300) and 75 Mbytes for hpuxcore (Series 800), plus space for additional applications. For a mixed cluster, you need space for both Series 300 and Series 800 filesets, plus space for additional applications. A mixed cluster server, for example, probably will need at least 400 Mbytes of space—plus space for future growth.

  You might want to mount a separate disk to contain the `/netdist` directory, from which the netdist server distributes files (for information on mounting files, see "Managing the File System" in this manual).

- For example, to distribute software for a complete Series 300 HP-UX system, you might mount a 7935 disk drive at `/netdist`.
- To distribute software for both Series 300's and Series 800's, the server must have files for both types of systems; a 7937 might be more appropriate for the `/netdist` directory.

- Update the server to HP-UX 7.0 (the `update` program is described in preceding modules of this chapter). Include the 7.0 LANLINK software. If you plan to set up anonymous ftp for file transfer (see below), also include ARPA.
- After updating to HP-UX 7.0, configure the 7.0 LANLINK software according to its documentation.
- To transfer files, you might want to set up anonymous ftp (file transfer protocol). Anonymous ftp is not the only method you can use to transfer files, but it provides a fast, effective vehicle for file transfer among networked systems. If you do not have anonymous ftp on your system, you can set it up using *sam*(1M) or by consulting the manual *Installing and Maintaining NS-ARPA Services.*

The procedure for configuring a netdist server is on the following page....

How to Configure a Netdist Server

1. Create the directory `/netdist` if it does not already exist on the server.
2. Make sure the `/etc/services` file contains an entry for the netdist server. The entry looks something like this:

   ```
   netdist          2106/tcp         # network file distribution
   ```

   If the entry does not exist, add it by typing the above example into the `/etc/services` file (use `vi` or any appropriate text editor).
3. Run `/etc/updist`. The `updist` program, a variant of `update`, transfers filesets from tape (or another server) to a hierarchical structure beneath the `/netdist` directory. These filesets are the ones the netdist server delivers to other systems on request: the filesets you select using `updist` are the only ones the netdist server can provide.

   To distribute software for both Series 300 and Series 800 systems, run the `updist` program with each set of 7.0 update tapes. After both sets of software are loaded, they are contained in separate directories beneath the `/netdist` directory.
4. When `updist` completes, type `/etc/netdistd -l` (lowercase `L`) to bring up the server. The `netdistd` program is the daemon that monitors the network for netdist requests; the `-l` argument causes the netdist daemon to log activity and error information to `/usr/adm/netdist.log`. The `netdist.log` file is described later in this module. For other options to `netdistd`, see *netdistd*(1M) in the *HP-UX Reference*.
5. Make the `update` program available to clients (see below).

Making the Update Program Available to Clients

After you've configured the netdist server, you must make the `update` program available to clients for network distribution.

1. Build the TOOL archive from the `/netdist` directory. To do this, type the following commands at the system prompt:
   - *Series 300.* To distribute software for a Series 300, type:

     ```
     cd /netdist/300/TOOL/product
     tar -cvf /tmp/TOOL.300 etc system
     ```

- *Series 800.* To distribute software for a Series 800, type:

```
cd /netdist/800/TOOL/product
tar -cvf /tmp/TOOL.800 etc system
```

- *Series 300 and Series 800.* To distribute software for the Series 300 and Series 800, type both sets of commands, as shown above.

2. This step assumes that you have anonymous ftp set up on your system (see "Prerequisites and Conditions" earlier in this module):

   Make the TOOL archive available to clients by putting the archives in a place where other systems can get them. Type one of these sets of commands at the system prompt:

   Series 300:
   ```
   mv /tmp/TOOL.300 /users/ftp/dist/TOOL.300
   chmod 444 TOOL.300
   ```

   Series 800:
   ```
   mv /tmp/TOOL.800 /users/ftp/dist/TOOL.800
   chmod 444 TOOL.800
   ```

Getting More Information

- For additional information on the `netdist` daemon, see *netdistd*(1M) in the *HP-UX Reference*.
- For additional information on using or setting up anonymous ftp, see *sam*(1M) or *Installing and Maintaining NS-ARPA Services.*
- For additional information on disk space, see the *HP-UX System Administration Concepts Manual.*
- For additional information on networking hardware, see available hardware configuration manuals and data sheets, such as the *HP9000 Series 300 Hardware Configuration Guide.*

Checking the netdist.log File

The `netdistd` daemon records events in `/usr/adm/netdist.log` if you specify the `-l` argument when you started the daemon (`/etc/netdistd -l`).

The `netdist.log` file contains the following information:

- The process ID of the parent `netdistd` daemon.
- The number of remote, incoming requests for the `update` program.
- The number of `update` program requests currently being serviced. By default, a netdist server can service twenty remote update sessions simultaneously. To change the default, use the `-C` option, as described in *netdistd*(1M) in the *HP-UX Reference*.

Here is an example of a `netdistd.log` file (the example is explained on the facing page):

`Version @(#)$Revision: 64.7 $`	*Startup*
`Building database 14:56:42`	
`Database ready 14:59:46`	
`netdistd.560 14:59:46 ... Started on port 2106`	*Parent daemon*
`netdistd.560 14:59:46 ... Connection limit is 2`	
`1.562 15:09:22 ... Begin service for hptest on port 2106`	*First child daemon*
`1.562 15:09:22 update @(#) $Revision: 64.492 $`	
`1.562 15:09:22 HP-UX hptest A.B7.00 B 9000/800 15453`	
`2.570 15:10:36 ... Begin service for hpclient on port 2106`	*Second child daemon*
`2.570 15:10:36 update @(#) $Revision: 64.492 $`	
`2.570 15:10:36 HP-UX hpclient 6.5 B 9000/300`	
`1.562 15:13:02 ... End service for hptest`	
`2.570 15:27:17 ... End service for hpclient`	

Startup: The first three lines print when you start the `netdistd` program. The first line shows the version of the `netdistd` program; the second and third lines indicate that the database was initialized successfully.

Parent daemon: Entries beginning with `netdist` indicate a "parent" daemon. The parent daemon's process ID (PID) is the number immediately following `netdist` (in the example, opposite, this number is `560`). When shutting down the netdist server, use this PID to kill the parent daemon process.

Child daemons: Entries beginning with a digit indicate "child" daemons. The first number (for example, the `1` in `1.562`) is a counter that increments by one each time a child daemon is spawned. This counter lets you determine the number of update sessions that have used this netdist server since the server was started. In the opposite example, two child daemons (`1.562` and `2.570`) have been spawned from the parent daemon (`560`).

The number following the counter (for example, the `562` in `1.562`) is the child daemon's process ID (PID). *Do not kill the child process*—if you do, you will interrupt an update in progress.

Shutting Down the Netdist Server

Prerequisites and Conditions

- The `netdistd` daemon spawns a child process to handle each incoming update request. Thus, if three update requests are being serviced, four copies of the `netdistd` daemon are running. *Do not kill the child process*—if you do, you will interrupt an update in progress.

Procedure

1. Determine which process is the parent process by viewing the file in `/usr/adm/netdist.log` (see "Checking the netdist.log File" for additional information on the `netdist.log` file.)
2. Kill the original parent daemon (see *kill*(1) in the *HP-UX Reference*), and wait for any child daemons to terminate normally. Killing the parent daemon prevents any new connections from being established.

Adding and Updating Filesets on the Netdist Server

Prerequisites and Conditions

- Read the previous module, "Shutting Down the Netdist Server."

Adding Filesets to an Existing Netdist Server

1. Run the `updist` program, as described in "How to Configure a Netdist Server."
2. When the `updist` program has completed take one of two actions:
   - Kill and restart the parent `netdistd` daemon (see "Shutting Down the Netdist Server" in this chapter).
   - Or, use the `kill` command to send the `SIGHUP` signal to the parent `netdistd` daemon (see *kill*(1) in the *HP-UX Reference*). This action causes the `netdist` daemon to re-read the `MAIN.pkg` file, which now contains the entries for your newly-added software (the `MAIN.pkg` file is described in the following module).

Updating Filesets on the Netdist Server

1. Shut down the netdist server (see "Shutting Down the Netdist Server" in this chapter). If you don't shut down the server, you might update files of a fileset that is being distributed to a remote system.
2. Run the `updist` program, as described in "How to Configure a Netdist Server."
3. When `updist` completes, restart the `netdistd` daemon by typing

   `/etc/netdistd -l` *(lowercase* L*)*

   The `-l` argument causes the netdist daemon to log activity and error information to `/usr/adm/netdist.log`. For additional information on the `netdist.log` file, see "Checking the netdist.log File" in this chapter.

Checking the MAIN.pkg File

After running the `updist` program, you might want to check the `MAIN.pkg` file, located beneath the `/netdist` directory. This file, which is automatically created or modified when you run the `updist` program, describes the filesets available for network distribution. The contents will look something like this:

```
source "/netdist/300/ACCOUNTING/netdist.pkg";
source "/netdist/300/ALLBASE1/netdist.pkg";
```

Statements beginning with `source` identify which filesets are available for network distribution. To prevent a fileset from being distributed, place a pound sign (#) at the beginning of the statement, or remove the statement from the file (see the opposite page for the complete procedure).

Dependencies exist between some filesets. A fileset that depends on one or more other filesets will have a comment directly beneath its `source` statement. For example:

```
 ...
source "/netdist/800/NS_SERV/netdist.pkg";
#NOTE:  NS_SERV depends on LANLINK
 ...
source "/netdist/800/NFS_RUN/netdist.pkg";
#NOTE:  NFS_RUN depends on LANLINK
 ...
source "/netdist/800/LANLINK/netdist.pkg";
 ...
```

In the example above, the NS_SERV and NFS_RUN filesets depend on the LANLINK fileset. Do not comment out the source statement for LANLINK unless you also comment out the lines for NS_SERV and NFS_RUN.

The `update` program issues an error message if the netdist server from which the program is updating contains unsatisfied fileset dependencies. For example:

```
ERROR:  Fileset "NS_SERV" depends on fileset "LANLINK", which does
not appear to be on this media.
```

To verify that the dependencies are correct, type the following command once the netdist server is running:

`update -s` *server_hostname* `-C`

Modifying the MAIN.pkg File

To modify the `MAIN.pkg` file once **`netdistd`** is running, follow these steps:

1. Shut down the netdist server (see "Shutting Down the Netdist Server" in this chapter).
2. Determine which filesets you do not want to distribute over the network.
3. Edit the `MAIN.pkg` file. Use `vi` (or any appropriate text editor) to place a pound sign (#) at the beginning of the statement or to remove the line. Be aware of possible fileset dependencies, as described on the facing page.
4. Re-start the netdist server with **`/etc/netdistd -l`** (lowercase `L`).

4

Starting and Stopping HP-UX

Starting and stopping HP-UX are routine, but critical, tasks. When the system is turned off, you can start up HP-UX in two ways. To stop a system, you must use an appropriate shutdown process. Simply turning the system off can corrupt the file system. When you change the system to an administrative state, you can reboot (restart) the system without turning it off; or you can shut the system down completely.

The following table shows the modules that describe these processes.

Module Name	How the Module Helps You
"Starting HP-UX"	Describes startup in attended and unattended modes, restarting HP-UX, interpreting messages, and coordinating startup with turning on devices (peripherals).
"Setting the System Run Level"	Describes run-levels, the system state, and changing the current run-level to a new one.
"Stopping HP-UX"	Explains the situations and procedures for shutting down the system.

You can get additional information about these processes in the *HP-UX System Administration Concepts* manual.

Starting HP-UX

You start up HP-UX when the system has been shut down completely (you turned the computer, and possibly an expander, off). When you partially shut down a running system to do administration, you can reboot the system without turning the computer off; or you can shut down the system completely.

Prerequisites and Conditions

- Some SAM tasks may restart (reboot) the system for you (for example, if you rebuild the kernel).
- To start up, your HP-UX system must have configured and installed hardware and software. The chapters named "Introduction to System Administration", "Constructing an HP-UX System", and "Updating HP-UX" have information about this.
- Start up an HP-UX cluster server as you would a standalone system. The chapter named "Managing an HP-UX Cluster" has information.
- The disk that contains the HP-UX file system can contain alternate HP-UX systems and other operating systems. If you want to boot a certain HP-UX system automatically, it must be the first system found by the bootROM.

- Your system must have certain files to start up properly (for example, `/etc/init`, `/etc/inittab`, `/dev/console`, `/etc/rc`). Without these files, the startup process will fail.
- The startup process may check the file system. This delays the startup, and you may need to perform additional tasks.
- If your system will not boot, you can use your recovery system to get a partial system going. The chapter named "Constructing an HP-UX System" has information.

The Startup Process May Check the File System

During the startup process, the system executes `/etc/fsclean`. This command determines the shutdown status of the system and returns three possibilities:

1. If the file systems were shut down properly, the startup process continues and you see the following message:

   ```
   /etc/fsclean: /dev/dsk/0s0 (root device) ok
   file system is OK, not running fsck
   ```

2. If any file systems were not shut down properly, the startup process is interrupted and you see:

   ```
   /etc/fsclean: /dev/dsk/0s0 not ok
   run fsck
   FILE SYSTEM(S) NOT PROPERLY SHUTDOWN,
   BEGINNING FILE SYSTEM REPAIR.
   ```

   At this point, the system runs `/etc/fsck` in a mode that can correct certain inconsistencies in the file systems without your intervention and without removing data. The `fsck` command will either:

   a. repair and reboot the system, incorporating the changes, or

   b. you may be asked to run `fsck` manually. If you need to run `fsck` manually, see the chapter named "Managing the File System".

3. If `fsclean` detects any other errors (for example, not being able to open a specified device file), you get an error message. The startup process can end, and you will need to solve the problem. The *Troubleshooting HP-UX* has information about possible problems in making HP-UX function.

Manual Startup Procedures

This section describes the startup procedures according to your purpose. Use the one you want. The *HP-UX System Administration Concepts* manual has detailed information about the startup process.

Starting a System in Unattended Mode

If the HP-UX system you want is the first system found by the bootROM and the system was shut down completely, you can start up that system in the unattended mode as follows:

1. Turn on all devices (peripherals) you want to use. Wait until they get into a ready state. You must turn on the disk that contains HP-UX.
2. You have two possibilities for turning on the computer:
   a. If you have only a computer (no expander), turn it on and go to the next step.
   b. If your computer is attached to an expander, proceed as follows:
      i. The CPU and the user-interface card (the card having the keyboard and such) should be in the computer, not in the expander.
      ii. In any case, especially if the cards are not installed as described above, turn on the expander. Then, turn on the computer and go to the next step.
3. The bootROM initiates startup and takes certain actions. Then, the HP-UX operating system takes control and completes the process. Watch the startup messages. Compare what starts up with what you expect to note possible problems. The exact messages depend on your configuration. The next two pages show hypothetical possibilities.
4. The startup process ends when you see the login prompt. If you do not get the prompt, the system did not start up. You will need to determine why. During the startup process, the system will perform a file system consistency check of the root disk if the system was shut down improperly. If your system is spread over multiple disks, you should perform a consistency check on the other file systems according to procedures described in the chapter named "Managing the File System".

Initial Startup Messages

```
Copyright
Hewlett-Packard Company.
All Rights Reserved.

BOOTROM Rev. C
Bit Mapped Display
MC68020 Processor
MC68881 Coprocessor
Keyboard
HP-IB
RAM 8388384 Bytes
HP98644 (RS232) at 9
HP98625 (HPIB) at 14
HP98643 (LAN) at 21, 0800009025C31


SEARCHING FOR A SYSTEM (RETURN To Pause)
RESET To Power-Up
```

After a short pause, the startup process continues, indicating that HP-UX is booting.

The next page shows the continuing messages.

HP-UX System Startup Messages in Unattended Mode

```
CONSOLE is ITE
ITE + 1 port(s)
MC68020 processor
Internal HP-IB Interface - System controller at select code 7
HP98644 RS-232 Serial Interface at select code 9
HP98625B High-Speed HP-IB Interface - system controller set at select code 14
HP98643 LAN/300 Link at select code 21
Bit Mapped Display at 0x560000
HP98620B DMA
real mem = 8376320
using 166 buffers containing 679936 bytes of memory
Local Link Address = (hex) 0800009003626
Root device major is 0, minor is 0xe0000
Swap device table: (start & size give in 512-byte units)
   entry 0 - autoconfigured on root device; start=2840, size=127976
-- BATTERY BACKED REAL TIME CLOCK
avail mem=3624960
lockable mem=5436780

<copyright information>

/etc/bcheckrc:
/etc/fsclean: /dev/dsk/0s0 (root device) ok
File system is OK, not running fsck
/etc/brc:
/etc/rc:
Is the date Fri Oct 27 10:47:01 PST 1989 correct? (y or n, default: y)
 ...
```

If you act in a few seconds, typing n lets you enter a new date. Otherwise, you get the default and the startup process completes. The chapter named "Constructing an HP-UX System" has information about setting the system clock.

Starting a System from Attended Mode

If you do not want to boot the first system found by the bootROM, you can start up HP-UX from the attended mode.

1. Turn on all devices (peripherals) you want to use. Wait until they get into a ready state. You must turn on the disk that contains HP-UX.
2. You have two possibilities for turning on the computer:
   a. If you have only a computer (no expander), turn it on and hold down the spacebar until you see `Keyboard` on the left side of the display. Wait until a list of available systems appears on the right side of the display. Go to the next step.
   b. If your computer is attached to an expander, proceed as follows:
      i. The CPU and the user-interface card (the card having the keyboard and such) should be in the computer, not in the expander.
      ii. In any case, especially if the cards are not installed as described above, turn on the expander. Then, turn on the computer and hold down the spacebar until you see `Keyboard` on the left side of the display. Wait until a list of available systems appears on the right side of the display. Go to the next step.
3. In the attended mode, the startup process pauses to show a list of available systems. You can see more than one system (even non-HP-UX systems). The hypothetical situation on the following page illustrates the idea of selecting a system from the attended mode:

`:HP7937, 1400, 0, 0`
`1H SYSHPUX`
`1D SYSDEBUG`
`1B SYSBCKUP`
	1H is probably your main HP-UX system. 1D is the debugger for the main system. 1B is the backup for the main system.

`:HP9144, 700, 1, 0` `2H SYSHPUX`	*An HP-UX system on a cartridge tape in an HP9144 drive, labeled 2H*

`:LAN, 21, hpfcma`
`3H SYSHPUX`
`3D SYSDEBUG`
`3B SYSBCKUP`
	3H is an HP-UX system that is available via your Local Area Network (probably a client in an HP-UX cluster). 3D is the debugger. 3B is the backup system.

Use the label to select the system you want to boot. For example, typing `2H` (you do not need (Return)) starts up the system on the cartridge tape in the HP 9144 drive that lets you install HP-UX.

4. Once you select an HP-UX system other than one shown for an installation tape, the startup process is the same as the process described for the unattended mode. You get the same messages that were shown earlier, and the process ends when you get the login prompt.

Changing System Run Levels

A run-level is a system state in which valid users can run a specific set of processes. Most of the time, the system functions in run-level **2**, the initdefault run-level that lets users log in and use system resources. On occasion, you change to run-level **s**, a special administrative or single-user state used for system administration tasks.

Prerequisites and Conditions

- Only the system administrator should change the run-level. Do not leave the system console unattended in this regard.
- Anyone who has write permission to `/etc/inittab` can create a new run-level or redefine the existing run-levels. In a listing of `/etc/inittab`, the permissions should look like this:

  ```
  -rw-r--r--  1 root  root  date information  /etc/inittab
  ```

  If the file does not have these permissions, execute:

  ```
  chmod 644 /etc/inittab
  ```

- The *System Administration Concepts* manual has detailed information about run-levels. Here is a brief look at them:

  0 A special run-level used for installation. Do not use this run-level.

  1 A single-user mode used during installation and maintenance. The system may reboot in single-user mode during some system administration tasks.

  s or S A special single-user (administrative) mode used for checking the file system or performing system backups. At run-level `s`, the `root` user can access the system from the system console. The only processes that can run on the system are the shell and the processes the `root` user invokes. Run level `s` is not recommended for routine use because processes that monitor the system do not execute.

2 A multi-user run-level that constitutes the system's normal operating state (the default run-level when the system starts up).

3 to 6 One can define (use) these run-levels for specific applications or commands. For example, the *X Environment* program uses run-level 3. By using this run-level, you get special login features within an X11 window. You can use the traditional login program by changing the system to run-level 2.

- Some system administration tasks require being in the single-user run-level (for example, performing a file system consistency check). In general, anytime you perform a major operation on a file system, rebuild the kernel, or otherwise make a significant change in HP-UX, determine if you need to be in the single-user run-level to perform the task. This manual tries to indicate each task that must be performed in this run-level. But it is possible to perform an undescribed task and need to be in this run-level.
- If you change to a different run-level, the processes corresponding to entries in `/etc/inittab` that do not accommodate the new run-level terminate. For example, if you have not added a respawn `getty` entry for run-level 3 for the console, entering run-level 3 from run-level 2 causes the console to die.
- Each terminal or RS-232 port used as an incoming terminal device must have a `getty` process executing. That is, for each incoming port, `/etc/inittab` must have a `getty` entry. For example, to accommodate a 3rd terminal in run-levels 2, 3, and 6 at 9600 baud, you could enter a line like this:

```
03:236:respawn:/etc/getty tty03 9600
```

  See *getty*(1M) in the *HP-UX Reference* manual for more information.

Procedures for Changing the Run Level

Changing to the Single-user Run Level

1. Tell users you are about to shut down the system. Have them save their work and then log off. Indicate you will provide a graceperiod of 30 seconds (for example) for them to log off when the shutdown process begins.
2. From the system console, log in as `root` and execute:

   ```
   shutdown 30
   ```

   The `/etc/shutdown` command terminates all processes except those that can run in the single-user state. During the shutdown process, the system warns users that they have a certain period to log off before the system stops. Not specifying a graceperiod gives users 60 seconds to log off.
3. When the shutdown process completes, use `ps -ef` to see which processes are still running. Any user-related process should not still be alive.

Changing to Run Level 2

When you change the system to a different run-level, you will probably want to return to run-level 2, the multi-user run-level.

1. Warn any users who may be logged in that you intend to change the run-level. Changing to another run-level while users are logged can kill their processes. You can do this verbally or use the `/etc/wall` (or `/etc/cwall` commands to broadcast a message that appears on the screens of all active users. See the chapter named "Constructing an HP-UX System" for information about these commands.
2. Become the root user.
3. Execute:

   ```
   init 2
   ```

4. Wait for the system to change to the new run-level.

Changing to a Numbered Run Level

To change to a numbered run-level:

1. Warn any users who may be logged in that you intend to change the run-level. Changing to another run-level while users are logged can kill their processes. You can do this verbally or use the `/etc/wall` (or `/etc/cwall` commands to broadcast a message that appears on the screens of all active users. See the chapter named "Constructing an HP-UX System" for information about these commands.
2. Log in as root.
3. After users log off the system, force the system to write the contents of its I/O buffers to disk by executing:

   ```
   sync
   sync
   ```

   The two `sync` commands ensure that all data is written to the buffers.
4. Use `/etc/init` to change the run-level (for example level 3):

   ```
   /etc/init 3
   ```

Defining New System Run Levels

You should ensure that people can run appropriate processes when the run-level changes. The `/etc/inittab` determines the available run-levels, so you edit this file to provide appropriate run-levels.

Prerequisites and Conditions

- Back up the original `/etc/inittab` file to a different name (for example, `/etc/orig_inittab`). If the new file does not work, you still have the original.
- The entries in `/etc/inittab` define how the system operates in each run-level. Each entry contains the following fields. Colons separate the fields.
  1. A one or two-character ID used to identify a process or process group (for example, `lp` for the line printer or `co` for the console).
  2. A list of run-levels to which each entry applies (for example, 236 for run-levels 2, 3, and 6).
  3. An action to be performed (for example, `initdefault`, `sysinit`, `bootwait`, `bootwait`, or `respawn`).
  4. The command that will be executed when that run-level is entered (for example, `/etc/bcheckrc` or `/etc/getty`).
- The *init*(1M) and *inittab*(4) entries in the *HP-UX Reference* have information about run-level entries.
- The *HP-UX System Administration Concepts* manual has detailed information about `/etc/inittab`.

Procedures

There is no exact procedure for defining run-levels. Instead, ensure that you provide for processes in the appropriate run-levels. The following items mention some things you should do. Then, you see two examples of `/etc/inittab` files. The idea is to show the process of determining the run-levels. By examining the items and the examples, you can extrapolate to your situation.

Making Sure the Console Works

You want the system console to have a `getty` (be able to log in) for every state you define. The following lines show possibilities:

```
co::respawn:/etc/getty console console
```

Recommended for the system console. Allows execution at all run-levels.

```
co:2345:respawn:/etc/getty console H #Don't use in single-user(PAM) state
```

Specifies run-levels 2 through 5

Providing for the Line Printer

You could have an entry such as:

```
lp::off:/bin/nohup /bin/sleep 99999999 < /dev/lp & stty 9600 < /dev/lp
```

Controlling the Startup Process

Certain entries in the `/etc/inittab` file determine the startup process. For example, the following line runs the bootrun command:

```
bc::bootwait:/etc/brc 1>/dev/syscon 2>&1 #bootrun command
```

Example /etc/inittab

The following example provides a system that has a system console and 3 terminals. Run-level `s` is the single-user run-level. Run-level 2 is a multiuser run-level, with a `getty` on every terminal. Run-level 3 is a test run-level, with a `getty` on both the system console and the system console (`/dev/tty01`).

```
is:2    :initdefault:
bl:     :bootwait:/etc/bcheckrc </dev/syscon >/dev/syscon 2>2&1
bc:     :bootwait:/etc/brc 1>/dev/syscon 2>&1
sl:     :wait:(rm -f /dev/syscon; ln /dev/systty /dev/syscon;) 1>/dev/console
rc:     :wait:/etc/rc <dev/syscon >/dev/syscon 2>&1
co:     :respawn:/etc/getty console  console
01:23   :respawn:/etc/getty tty01   9600
02:2    :respawn:/etc/getty tty02   9600
03:2    :respawn:/etc/getty tty03   9600
```

Here is another, customized, example that shows possibilities:

```
is:23:initdefault:
st::sysinit:stty 9600 clocal icanon echo opost onlcr ienqak ixon
icrnl ignpar < /dev/systty
#bootlog follows
bl::bootwait:/etc/bcheckrc </dev/syscon >/dev/syscon 2>&1
bc::bootwait:/etc/brc 1>/dev/syscon 2>&1 #bootrun command
cr::bootwait:/bin/cat /etc/copyright >/dev/syscon
lp::off:/bin/nohup /bin/sleep 999999999 < /dev/lp & stty 9600
< /dev/lp
#run com
rc::wait:/etc/rc </dev/syscon >/dev/syscon 2>&1
#power fail routines follow
pf::powerfail:/etc/powerfail 1>/dev/console 2>&1
#single_user environment
su:1:respawn:/etc/single_user 0</dev/syscon 1>/dev/syscon 2 >&1
#multi_user environment setup
mu:23:wait:/etc/multi_user 0</dev/syscon 1>/dev/syscon 2>&1
#Don't run following line in single_user(PAM) state
co:023456:respawn:/etc/getty console H
# SC,BA Patch Name
zz::off:                Line
t0:23:respawn:/etc/getty      -t120    tty10 2400 # 13,0
t1:23:respawn:/etc/getty               tty11 H    # 13,1
```

Shutting Down the System

You should never just turn an HP-UX system off! Instead, shut the system down properly. In an HP-UX cluster, you can turn a cnode for a client off, but do not turn the cluster server off. Typically, you shut down the system down for one of two reasons:

1. Get into the single-user state so you can do system administration tasks such as update the system, reconfigure the kernel, or check the file systems, or
2. Shut down the system totally to perform a task such as installing a new interface card.

Prerequisites and Conditions

- Stopping the system improperly can corrupt (damage) the file systems. Never stop the system by turning it off!
- Only the system administrator or a designated superuser should shut down the system.
- The `/etc/shutdown` command warns users of impending shutdown; halts daemons; kills unauthorized processes; unmounts file systems; puts the system in single user mode; and writes the contents of the I/O buffers to a disk. You see several messages during the process. You should watch them to note actions and possible problems.
- The `shutdown` command warns all users to log off the system, using a graceperiod you can specify. If you do not specify one, users get 60 seconds to log off. You should notify active users as to when the system will be shut down. Give them enough time to finish their work and log off. You can do this physically or use the `/etc/wall` or `/etc/cwall` commands. The chapter named "Constructing an HP-UX System" has information.

- In an HP-UX cluster, clients need only log off, but the shutdown process works better if you turn the cnodes off. Do not shut down the system from a client; do the work from the root server.
- If you use a network service, do not run `shutdown` from a remote system via `rlogin`. The shutdown process logs you out prematurely and returns control to the system console.
- You cannot use SAM to shut down the system, but some SAM tasks shut down to the single-user state, or they ask you to do so.
- See the *shutdown*(1M) entry in the *HP-UX Reference* manual for information about options and features.
- The *System Administration Concepts* manual has information on system shutdown concepts.

Manual Procedures

Going to the Single-user State for Maintenance

1. As the root user, change to the root directory if not already there:

   ```
   cd /
   ```

2. Shut down the system. You have some alternatives for doing this. Also, the shutdown process asks if you want to send a message. If you elect to broadcast a message, respond with y and then type the message. When you finish, press (Return) (or (Enter)), and then (CTRL)-(D). The following examples show alternatives for shutting down to the single-user state:

`shutdown`	*Shuts down to single-user state, allowing the default 60 second graceperiod*
`shutdown 0`	*Shuts down the system with no graceperiod*
`shutdown 30`	*Begins the shutdown to the single-user state after a 30-second graceperiod*

3. While the system is in the single-user state, perform the necessary system administration tasks. When you finish, you can start up the system without turning off anything by executing:

   ```
   reboot
   ```

   As always, watch the messages to see that everything is happening correctly.

   Some system administration tasks will do the rebooting for you.

Shutting Down the System Completely

1. As the root user, change to the root directory if not already there:

```
cd /
```

2. Shut down the system. You have some alternatives doing this. Also, the shutdown process asks if you want to send a message. If you elect to broadcast a message, respond with y and then type the message. When you finish, press [Return] (or [Enter]), and then [CTRL]-[d].

   From the multi-user state, you can shut down the system completely:

```
shutdown -h
```

   This process is rather harsh and sudden. It is generally better to take the system down in steps:

   a. Execute:

   `shutdown 20` *This gets the system into the single-user state, allowing a 20 second graceperiod*

   b. Execute:

   `reboot -h` *This brings the system to a complete stop*

   You know the system is shut down completely when the system displays `halted` and pressing a key takes no action.

3. When the system is halted, turn the system off as follows:

   a. If you have only a computer (no expander), turn the computer off. Then, turn the devices off as required.

   b. If you have a computer and an expander, turn the computer off, turn the expander off, and then turn the devices off as required.

4. When you want to start up the system again, see the earlier procedure for starting up HP-UX.

Shutting Down the System to Activate a New Kernel

You may want to shut down the system only to activate a new kernel. To do this, execute:

```
shutdown -r 0
```

The `-r` option causes the system to reboot immediately after the system gets into the single-user state.

Do not execute `shutdown -r` from run-level s. You must reboot using the `reboot` command.

5

Managing Groups and Users

To log in, a person must have one or more HP-UX accounts. The system administrator manages the accounts that let users access the HP-UX system. The system administrator also manages the groups that contain users.

Because of the sequencing of minor tasks within a procedure, you may encounter information about users and groups while you perform a certain task. When you work manually, manage the group and then manage the user. SAM does some tasks that let you manage the users and their groups at once. This chapter describes tasks related to groups first. Then, it describes tasks related to users.

Managing a Group

Each user must be a member of a group. If you manage a large number of local and remote users, and if several users need more than one account, this can become a complex task. Besides just creating groups for users, take time to plan the process and work systematically.

Prerequisites and Conditions

- A user must belong to some group. The entry for the root user (the superuser) in `/etc/group` should look like this:

  ```
  root::0:root
  ```

  It can look like the following entry if you want to allow users named `brick` and `prod` to function as the superuser.

  ```
  root::0:root,brick,prod
  ```

- A user can be in more than one group.
- Within a group, or across groups, a physical user can have more than one user account. For example, a John Quincy Smith could have user accounts named `johns`, `jqs`, or even `grp-ltr` within a group, or this person could have these accounts across several groups.
- Within a group, two physical users should not share the same user account unless that is an acceptable arrangement.
- The group should exist before you manage users. To add a user manually, the group must already exist. If you use SAM to add a new user, specifying the group for the user creates the group if it does not exist.
- The *group*(4) entry in the *HP-UX Reference* has information about groups.
- You should link `/etc/group` with `/etc/logingroup` to ensure that the two files contain the same information. Execute:

  ```
  ln /etc/group /etc/logingroup
  ```

SAM Procedure for Managing Groups

`System Administration Manager`

↓

`Groups ->`

↓

At this point, you have the following items:

- `Add a New Group to the System ...`
  Under this item, you complete the following fields:
  - Group name
  - Login names of group members
- `Remove a Group from the System ...`
  Under this item, you accommodate the following fields:
  - Group name (a pop-up)
  - Group to assign files to
  - Group members
- `View/Modify Group Membership ...`
  Under this item, you accommodate the following fields:
  - Group name (a pop-up)
  - Group members

Manual Procedure for Creating a Group

Add an entry to /etc/group. This file has lines like those that follow:

```
root::0:root
other::1:root,daemon,uucp,who,date,games,sync
bin::2:root,bin,daemon,lp
sys::3:root,bin,sys,adm
adm::4:root,adm,daemon
daemon::5:notes,anon,uucp
mail::6:root
rje::8:rje,shqer
guest::10:guest,emg,jem,spider
news::12:news
users::20:john,archive,sue,richard,,johnt,krm,jaci
prod::30:maryb,judyg,ginger,pams,marlene,stevem,prod
```

In the above example, colons separate the four fields as follows:

1. The name of the group.
2. An encrypted password or appropriate substitute.
   a. Insert an asterisk (*) to prevent members in other groups from being switched to the group.
   b. Leave the field blank to allow non-group members to be switched to the group.
3. The numerical group-ID number (GID). Each group must have a unique GID. Choose a convenient number, making sure it does not conflict with the number of an existing group in /etc/group. If some users already have a GID on a remote system, using the same number for their group on the local system lets them execute commands on the remote and local systems.
4. A list of names of users who belong to the group, separated by commas with no spaces. (You can let a long line wrap around on the display. Just type the entry as one line.)

Manual Procedure for Changing Group Membership

Use `/bin/chgrp` to change the membership field in `/etc/group`.

For example, to change the group ownership of the file `/users/acctinfo` to the `finance` group, execute:

```
chgrp finance /users/acctinfo
```

Providing Privileged Groups

Privileged groups control user access to HP-UX features such as real-time priorities, file locking, and changing file ownership.

Prerequisites and Conditions

- By default, your HP-UX system allows all users *CHOWN* (change ownership) privileges. HP recommends that you keep the system in this state. In most cases, you can accommodate system administration and users without giving certain groups a privileged status.
- The *setprivgrp*(1M) entry in the *HP-UX Reference* manual describes the features and capabilities of privileged groups.

There is no SAM procedure for providing privileged groups. The manual procedure appears on the following page.

Manual Procedure for Establishing a Privileged Group

The `/etc/setprivgrp` command changes group privileges.

`setprivgrp -g|-n|` *groupname* `[`*privileges*`]`
`setprivgrp -f` *filename*

The parameters work as follows:

`-g`	All groups.
`-n`	No groups.
groupname	The group whose privileges are being changed.
privileges	See the following list.

For *privileges,* you can use the following values:

RTPRIO	Real-time priorities
MLOCK	Locks processes into memory
CHOWN	Changes file ownership (default privilege)
LOCKRDONLY	Permits locking of read-only files using the `lockf(2)` system call.
SETRUGID	In existing code, permits the user to set only the real user ID and real group ID using the *setuid*(2) and *setgid*(2) system calls. For new code, use *setresuid*(2) and *setresgid*(2) to permit changes to the real user ID and real group ID of a process.

The system stores these privileges in memory. You must reset them when you reboot the system. You typically set privileges in `/etc/rc`. The `-f` option causes privileged groups to be set from the named file, which typically is `/etc/privgroup`.

Adding a User Account (Local or Remote)

A person cannot use HP-UX resources until you give that person an account within a group.

Prerequisites and Conditions

- Users relate to groups as follows:
  - Users fit into groups that have names such as `users`, `root`, `other`, and `mktg`.
  - Each user must belong to at least one group. A user can belong to several groups.
  - When you add a new user, you either:
    - add the person to an existing group, or
    - create a new group and then add the user to that group. If there is no existing group, see the later section named "Creating Groups". Then, return to this section to add the user.
- If you have networking that provides for a remote login (the NS-ARPA Services, for example), a remote user may request an account on your system. You can add this user as usual. Use the login name and user ID that the user has on the remote system.

The procedures begin on the following page.

SAM Procedure for Adding a User

`System Administration Manager`

↓

`Users ->`

↓

`Add a New User Account to the system ...`

Under this menu option, you accommodate for the following items:

- `Login name` (for example, `jtq`, `janek`)
- `Primary Group name` (for example, `users`, `other`)
- `Home directory` (for example, `/users/janek`)
- `Start-up program` (for example, `/bin/ksh`)
- `Login with X11 windows? (y or n)`
- `Real name` (for example, `Jane Kane`)
- `Office location` (for example, `1UG7`)
- `Office phone` (for example, `555-1234`)
- `Home phone` (for example, `555-4321`)
- `Password` (a pop-up) (entering `,..` forces a user to provide a password)

Manual Procedure for Adding a User

To give a user an account (a login), complete the following tasks:

1. Add the user to a group (for example, `users`). in `/etc/group` as follows:
   a. See the earlier module named "Managing a Group".
   b. Add the user, creating the group if necessary.
   c. Come back to this module to finish adding the user.
2. Add an entry to `/etc/passwd` for the new user. In this entry, the password field determines how a user obtains a password:
   a. Use `,..` in the password field to force a user to provide a personal password. Editing the `passwd` entry so the password field contains `,..` is the recommended way to add a user.
   b. You can leave the password field blank. Then, after you add a user, you can log in as `root`, execute `passwd` *user_name,* and provide a password as prompted by the system. HP does not recommend this scheme because it represents a potential security problem.
3. Create the user's home directory. Typically, you place users in a file system such as `/users` or `/others`. You may use a scheme in which you place users in a directory under `/users` (for example, `/users/prod` where `prod` might be people in a production department).
4. Create the user's login files or help the user create the files. This refers to the "dot" files that go in the user's home directory (for example, `.profile`, `.mailrc`, and `.x11start`). The earlier chapter named "Constructing an HP-UX System" has information about these files. You can also get information about them in the *HP-UX System Administration Concepts* manual and in *A Beginner's Guide to Using Shells.*

The next several sections have information about how to complete these tasks.

(1) Add the User to a Group

1. Examine the `/etc/group` file to see the list of existing groups.
2. If you have a group to which you add users, use an editor to include the new user. For example, existing groups for users named `users` and `mgmt` might look like this:

   ```
   users::20:jake,sue,kmh,jtb,archive
   mgmt::30:samk,janicew,kjl,budget
   ```

   To add a user whose login name is `maryt` to the group named `users`, edit the entry, including a comma and the login name after `archive`.
3. If no group for users already exists, you must create the group and then add the user according to the above steps. An earlier module named "Managing a Group" explained how to create a group.

(2) Add the User to the /etc/passwd File

The `/etc/passwd` file must have an entry for each user. If you have a secure system, the *HP-UX System Security* manual has information about modifying `/.secure/etc/passwd` as well as `/etc/passwd`. The *passwd*(4) entry in the *HP-UX Reference* has information about `/etc/passwd`.

1. Log in as `root`.
2. You may want to make a backup copy of `/etc/passwd` before you modify it:

   ```
   cp /etc/passwd /etc/passwd.old
   ```

3. Edit the `/etc/passwd` file to add a line for each new user as shown by the following example:

   ```
   louis:,..:110:300:Louis Smith,x4321:/users/louis:/bin/ksh
   ```

   The items on the following page explain the fields.
4. Using `/etc/pwck` when you finish editing checks for inconsistent entries. For example, entering `lsh` instead of `ksh` for the shell would display:

   ```
   louis:,..:110:300:Louis Smith,x4321:/users/louis:/bin/lsh
           Optional shell file not found
   ```

`louis` The login name (1-8 characters). The first must be alphabetic.

`,..` Forces the user to provide a password. The system encrypts the password supplied by the user so it looks like `RZ5MagE5Lhn31`.

`110` The user-ID (UID) number. The superuser has UID 0. If a remote user already has a UID, assigning the same number ensures that the user can function on the remote and local systems. Use 0 for `root`; the system reserves 1-100 and 200 and 201; use values between 100 and 60001; UIDs greater than 60001 map onto 60001.

`300` The group-ID (GID) number assigns the user to the group having that number.

`Louis Moore,x4321` A comment field for user information (full name, phone extension, and so on up to four subfields, separated by commas).

`/users/louis` The pathname to the user's home directory

`/bin/ksh` The shell that executes when the user logs in (default is `/bin/sh` for the Bourne shell).

(3) Create a Home Directory for the User

Use `/bin/mkdir` to create a home (login) directory for the user. The following example creates a home directory for a user named `louis` with appropriate permissions and ownership:

1. Go to the directory you specified for the user in `/etc/passwd`. For example:

   ```
   cd /users
   ```

2. Create a home directory for the user:

   ```
   mkdir louis
   ```

3. Set the ownership of the directory:

   ```
   chown louis louis
   ```

4. Set the group ownership:

   ```
   chgrp users louis
   ```

5. Set the permissions to allow only `louis` to write to the directory:

   ```
   chmod 755 louis
   ```

(4) Create a Login File for the User

The final step in manually adding a user is to provide a login file. The straightforward method is to copy a system login file to the user's home directory, renaming the file. The table shows the correspondences according to the user shell provided by the `/etc/passwd` entry, the chosen system login file, and the name given to the login file in the user's home directory.

For This Shell ...	Copy This System File ...	Rename It as Follows ...
Bourne	`/etc/d.profile`	`.profile`
Korn	`/etc/d.profile`	`.profile` Some Korn shell users set up a secondary login file named `.kshrc`. *A Beginner's Guide to Using Shells* has information.
C	`/etc/d.cshrc` `/etc/d.login`	`.cshrc` and `.login`

For example, for a Bourne shell user named `john`, who is in the group named `users`, you might copy a login file as follows:

```
cp /etc/d.profile /users/john/.profile
```

Set the owner, group, and permissions for the login file as shown earlier for the home directory.

If you add users often, you may want to create a set of customized environment files.

- The previous chapter named "Constructing an HP-UX System" describes many environment files.
- The *HP-UX System Administration Concepts* manual has information about environment files and variables.
- The *Beginner's Guide to Using Shells* has examples of login and environment files.

Viewing or Modifing User Account Information

You can use SAM to look at and change the information for a user (for example, the startup program or user group).

`System Administration Manager`
↓
`Users ->`
↓
`View/Modify a User's Account Information ...`

Under this option, you can examine or modify the following items:

- `Login name` (a pop-up)
- `Login name`
- `Primary Group name`
- `Home directory`
- `Start-up program`
- `Real name`
- `Office location`
- `Office phone`
- `Home phone`
- `Change user's password? (y or n)`

Manual Procedure for Viewing or Modifying Account Information

Use typical HP-UX commands (`more`, `cat`, `grep`, `vi` and so on) to examine or edit the files that contain the information shown above for the SAM procedure.

Removing a User Account

It is easy to accumulate users. On occasion, you should monitor the file system that contains users and determine if they still need an account.

Prerequisites and Conditions

- Do not simply remove users that are no longer physically in a group. Users may be logically connected to a group via electronic mail. They may have some other, unforeseen, connection with your system (for example, using your printer or plotter). A user may need to access your system via a remote login.
- Before you remove a user, see if the person still needs an account. If the account is not needed, determine the best way to transfer the person's files to a backup media or another system. The chapter named "Backing Up and Restoring the System" has information about backup methods. In most cases, using `find`, `cpio`, and `tcio` with pipes to back up the user's files to a cartridge tape works well.
- Your company may have specific policies concerning how you handle the files of a user whose account has been removed. Be sure to check into this.
- The *HP-UX System Security* manual has information about removing a user from a secure system.

SAM Procedure for Removing Users

`System Administration Manager`
↓
`Users ->`
↓
`Remove a User Account from the System ...`

Within this option, complete the following items:

- `Login name` (a pop-up)
- `Remove all files/directories? (y or n)`

  Within this item, you have the following options if you answer no to the previous question:

  - `Assign all files/directories to another user` If you say yes, a pop-up asks for the login name.
  - `Remove only home directory and its files`
- `Home directory` (You see this only when directory name will not fit on display)
- `Start-up program` (You see this only when program name will not fit on display)

Manual Procedure for Removing Users

1. Remove the user's home directory and any other personal directories. The following command string removes all the user's files and directories. Be sure that no other users need them. For example, for a user named `john`, execute:

   ```
   find / -user john -exec rm {} \;
   ```

2. Delete the user's entry (or entries) from the `/etc/passwd` and `/etc/group` files.

Deactivating a User Account

You need not always remove a user. You may want only to deactivate a user, probably while the user is away from the system for an extended time. Deactivating a user keeps the user from logging onto HP-UX, but otherwise leaves the user intact on the file system.

Prerequisites and Conditions

- Before you deactivate a user, it is a good idea to back up the user's files. Use an appropriate backup method. The chapter named "Backing Up and Restoring the System" has information about backup methods. In most cases, using `find`, `cpio`, and `tcio` with pipes to back up the user's files to a cartridge tape works well.
- You archive the tape in a safe place or give it to the user. This depends on the reasons for deactivating the user and any policies your company may have (for example, system security or dealing with sensitive material).

SAM Procedure for Deactivating Users

`System Administration Manager`

↓

`Users ->`

↓

`Deactivate a User Account ...`

Under this option, you respond to the following items:

- `Login name` (a pop-up)
- `Special processing on user's files (y or n)`
  Under this item, you have the following items if you answer no to the previous question:
  - `Assign all files/directories to another user`
  - `Remove only home directory and its files`
  - `Remove all files and directories?`
- `Home directory` (Appears only when directory name will not fit on display)
- `Start-up program` (Appears only when program name will not fit on display)

Manual Procedure for Deactivating Users

You deactivate a user's login account by placing an asterisk (*) in the password field in the person's entry in `/etc/passwd`.

For example, to deactivate the user named `louis`, change the entry in `/etc/passwd` as follows:

```
louis:*:110:300:Louis Moore,x443:/users/louis:/bin/ksh
```

Then, `louis` cannot access the system until you delete the asterisk and assign or force a new password.

Reactivating a User Account

If you deactivate a user, you can subsequently reactivate the user.

SAM Procedure for Reactivating a User

`System Administration Manager`

↓

`Users ->`

↓

`Reactivate a User Account ...`

At this point, you respond to the following items:

- `Login name` (a pop-up)

Manual Procedure for Reactivating a User

Edit the password field in the entry for the user in the `/etc/passwd` file. Remove the asterisk (*). Insert `,..` in the field to force the user to provide a new password during the initial login.

6

Managing the File System

You manage the HP-UX file system by exploring how it best serves the needs of users. Because the system is huge, and because it contains many types of "software" (commands, utilities, tools, applications, languages, environments, libraries), learning how to manage everything can take considerable time. The modules in this chapter can help you perform tasks, but eventually, you need to understand the file system to manage it effectively.

The following items orient you to the chapter:

1. A table on the following page orients you to managing the file system.
2. Following the table, a section identifies the major file systems used by system administrators.
3. Then, the task-modules begin.

Major File System Tasks

Task	Description	Can Use SAM?
Adding or Removing a Disk Drive	Use the manual for installing peripherals for your series, and possibly some tasks described in this manual	Yes, for hard disks.
Planning Media Initialization	Situations related to initializing media.	No
Initializing Media in HP-UX Format	Prepares media for holding data or a file system.	Hard Disks
Creating a LIF Volume	Relates to media for BASIC and Pascal systems.	No
Creating a File System	Prepares media for holding files (directories).	Yes
Mounting/Unmounting a File System	Associated a file system with a device.	Yes
Moving a File System onto Another Disk	Distributes an HP-UX system over multiple disks. Also explains how to manually mount a file system.	No
Automatically Mounting and Unmounting File Systems	Shows how placing entries in `/etc/checklist` has the system automatically mount file systems. It also explains how they get unmounted.	Yes

Task	Description	Can Use SAM?
Checking File System Consistency	Repairing corrupt media. The process can become complex.	No
Adding or Deleting Swap Space	Increasing the swap space set up during installation.	Partially
Deleting Swap Space	Reducing swap space, which frees some disk space.	Yes
Adding Dynamic Swap to a File System	Provides temporary relief for insufficient swap space.	Yes
Transferring Files	Transferring file according to systems and situations.	Certain Situations
Automating Processes	Explains how to use `cron` to perform tasks periodically.	No
Monitoring the File System	Explains the use of log files.	No
Enabling Long Filenames	Provides for specifying long file names.	Yes
Removing a Product from HP-UX	Explains how to remove the filesets that constitute a product.	Some Products
Adding/Removing an NFS File System	Relates to Yellow Pages and such.	Yes

File Systems Used by System Administrators

The *HP-UX System Administration Concepts* manual has information about the file systems. This section acquaints you with the file systems and directory structure used by system administrators.

The following figure shows the major file systems you use and should understand. Tables on following pages describe the file systems.

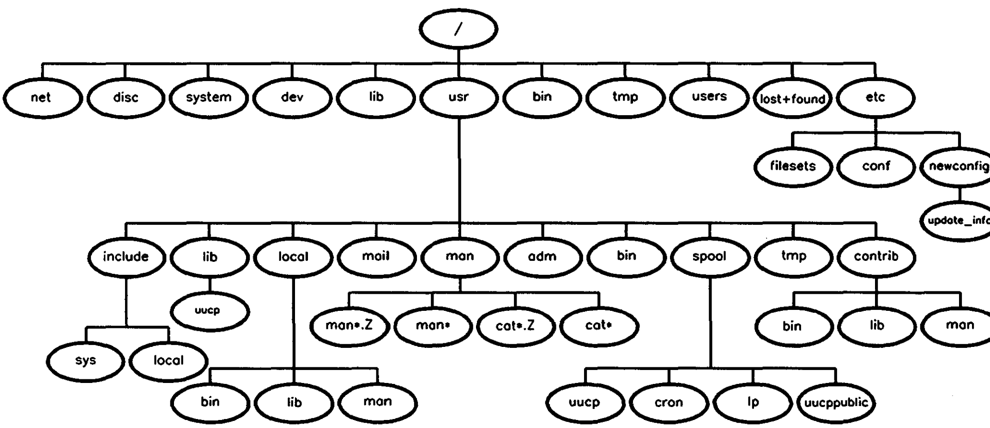

Tables that explain the files begin on the next page.

Directory	Description
/bin	Contains compiled, often-used commands.
/dev	Contains block and character special device files used to communicate with devices. See *mknod*(1m).
/disc	An empty directory created by the installation; used temporarily for mounting other file systems.
/etc	Contains most system administrator commands and configuration (customization) files.
/etc/newconfig	Receives customized configuration files and shell scripts during an update so you can use them for reference. You typically copy many of these files back into `/etc`.
/etc/newconfig/ Update_info	Contains release information. During an update, read these files before you start the update program.
/etc/conf	Contains kernel configuration description files.
/etc/filesets	A list of loaded filesets (do not delete anything).
/lib	Contains object code libraries and related utilities.
/system	Contains revision lists and customize scripts from installations and updates (do not delete anything).
/tmp	Used to hold files temporarily.
/users	Contains directories for users and the files and subdirectories they create.
/usr	Contains commands, log files, and so on.

The explanations continue on the next page.

Directory	Description
/usr/adm	Contains system administration data files.
/usr/bin	Contains commands not required to boot, restore, or repair the file system.
/usr/contrib	Contains files and commands contributed by user groups.
/usr/contrib/bin	Contains contributed commands.
/usr/contrib/man	Contains online documentation for contributed commands.
/usr/include	Contains high-level C language header files; the shared definitions.
/usr/include/local	Contains site-specific C language header files.
/usr/include/sys	Contains low-level, kernel-related C language header files.
/usr/lib	Contains less-used object-code libraries, utilities, lp commands, and miscellaneous data files.
/usr/local	Contains localized, site-specific files.
/usr/local/bin	Contains localized, site-specific commands.
/usr/local/lib	Contains object code libraries for the site-specific commands.
/usr/local/man	Contains online documentation for the site-specific commands.

The explanations continue on the next page.

Directory	Description
/lost+found	Created by `newfs` and used by `fsck` for orphaned files and directories.
/usr/mail	Used by the mail facilities for your mail box.
/usr/spool	Receives spooled (queued) files for various programs.
/usr/spool/cron	Receives spooled jobs for `cron` and `at`.
/usr/spool/lp	Contains control and working files for the `lp` spooler.
/usr/spool/uucp	Receives queued work files and contains lock files, log files, status files, and other files for `uucp`.
/usr/spool/uucppublic	Used for free access of files to remote systems via LAN and `uucp`.
/usr/tmp	An alternative directory for placing temporary files, typically very large files.
/usr/man/cat1 ... cat9	Contains online documentation that has already processed to speed up access.
/usr/man/cat1.Z ... cat9.Z	Compressed versions of `cat` directories.
/usr/man/man1 ... man9	The unformatted online documentation pages.
/usr/man/cat1.Z ... cat9.Z	Compressed versions of the online documentation pages.

Adding or Removing a Disk Drive

After you have an installed or updated an HP-UX system, you may want to add another disk drive, or remove it later.

Prerequisites and Conditions

- You cannot have more than 8 disk drives on an HP-IB port.
- Cartridge and 9-track drives typically use an HP-IB port at Select Code 7.
- Flexible disk drives typically use an HP-IB port at Select Code 7.
- Hard disk drives typically use a high-speed HP-IB port at Select Code 14.
- Optical disks require a SCSI adapter at Select Code 47.
- SAM accommodates only hard disk drives.

SAM Procedure for Adding a Hard Disk Drive

```
System Administration Manager
      ↓
File System ->
      ↓
Add a Hard Disk Drive ...
```

At this point, you see the data-entry screen. Enter appropriate data, realizing that you can also perform tasks such as initializing the media and creating a file system. If you need help with a particular field, call the Help Screen for that field.

Manual Procedure for Adding Any Disk Drive

Find the procedure and information in the manual for installing peripherals for your series.

SAM Procedure for Removing a Hard Disk Drive

`System Administration Manager`
↓
`File Systems ->`
↓
`Remove a Hard Disk Drive ...`

At this point, select a disk to remove. Enter appropriate data in each field. If you need help with a particular field, press the `Help` softkey.

Manual Procedure for Removing Any Disk

1. If appropriate and desired, backup files on the disk to be removed.
2. Remove the mount directory and device file.
3. Unmount the disk and remove its entry from `/etc/checklist`. If the disk is a swap device, you may need to reboot HP-UX.
4. Turn the device off and physically remove it.

Planning Media Initialization

This module discusses initializing storage media so you can then select an appropriate procedure. As guidelines, you must initialize new optical and flexible disks. You should initialize new hard disks. HP recommends initializing new cartridge tapes, but this is not typically necessary. You should initialize any media you suspect is corrupted and contains no valuable data.

Uses of Storage Media

To Hold an HP-UX File System

Use a hard disk. You can use an optical disk when you want to backup a file system and subsequently be able to mount the file system to facilitate restoring files. You can use a flexible disk, but you have limited capacity. Do not use cartridge or magnetic tapes. They have slow access and the tape has a short life span due to wear and tear.

To Store or Transfer Files

- Use magnetic 9-track tape or optical disks when you need large capacity (for example, an HP-UX cluster that has about 600 Mbytes).
- For standalone workstations, you can use the above media and cartridge tapes. Cartridge tapes have 16 or 67.5 Mbyte capacities for 150-foot and 600-foot tapes.
- You can use flexible disks, but they seldom have sufficient capacity (about 1 Mbyte).

To Transfer Files to Systems with LIF Format

Use cartridge tapes or flexible disks or cartridge tapes having files you changed into Logical Interchange Format (LIF) volumes to transfer files between the HP-UX operating system and the HP BASIC and Pascal systems.

Unformatted Media

You may want only to store data or files (for example, streamed data or backup files). You typically use optical disks or cartridge tapes. Some procedures do not allow the use of flexible disks, and in most cases, they have limited capacity. Use `/usr/bin/mediainit` for the initialization. Later, use commands such as `cpio`, `tcio`, `find`, `tar`, `fbackup`, and `dd` to store the data or files.

Alternative Media

HP-UX has software that lets you use files with an MS-DOS ™ format. Use `/usr/bin/mediainit` for media containing HP-UX files when you run the HP-UX operating system. Use `FORMAT` (an MS-DOS command) for media containing MS-DOS files when you run the MS-DOS Operating System (See the *MS-DOS Reference* manual for information).

Later, you can use `dos2ux` and `ux2dos` to format files for use on HP-UX or MS-DOS. See the *HP-UX Reference* manuals for information about these commands.

Getting Information about Drives

To get information about a particular drive (for example, installation procedure, select code, bus address, device file names, device drivers, and so on), see the manual for installing peripherals for your series.

Getting to Actual Initialization

After you decide how you want to use media, move ahead to the appropriate module to initialize the media.

Initializing Media in HP-UX Format

Initialization prepares the media for storing files by setting up such things as cylinder groups, cylinders per group, tracks, sectors per track, bytes per sector, interleave factor, and so on. After initialization, the media can contain data, files, or a file system.

Prerequisites and Conditions

- The drives containing the media to initialize must have raw device files in `/dev` or in a directory such as `/dev/rdsk`.
- The kernel must have a device driver for each class of drives. The file named `/etc/conf/dfile` shows the active device drivers.
- Do not initialize media that contains a mounted file system. The initialization destroys all existing data.
- In general, you should be in the single-user state to initialize hard disks. You need not be in the single-user state to initialize optical disks, 9-track tapes, cartridge tapes, or flexible disks if you make sure users cannot access the media during initialization.
- Knowing the relationships among drives, device files, and device drivers helps you avoid inadvertently initializing the wrong media. For example, using `/dev/dsk/3s0` as the block device file for a 3rd hard disk provides a useful correspondence.
- Do not change defaults for initializing media unless you know why you want to change them. For more information on this, see the *HP-UX System Administration Concepts* manual.

SAM Procedure for Initializing HP-UX Media

You can only initialize media as you add a hard disk.

```
System Administration Manager
     ↓
File Systems ->
     ↓
Add a Hard Disk Drive ...
```

At this point, you can enter data related to initializing a hard disk under the question about creating a new file system. If you need help, call the Help Screen.

Manual Procedure for Initializing HP-UX Media

1. Log in as the root user. While you need not be in the single-user run level, you must ensure that no one accesses the media during initialization.
2. If necessary, turn the drive on. Wait until it completes self-tests and runs steadily.
3. If necessary, insert the media and wait until it is ready for use.
4. Use `/usr/bin/bdf` to make sure the drive is not mounted (the device does not appear in the listing).
5. Initialize the media using the character special device file.

   `mediainit` [*options*] [*path_and_char_device*]

   You seldom need options except, perhaps, to specify an interleave factor. The *HP-UX Reference* manual has more information. The following examples show typical instances:

`mediainit /dev/rdsk/4s0`	*Initializes a hard disk drive set to bus address 4*
`mediainit /dev/rscsi`	*Initializes an optical disc in a drive connected to a SCSI adapter*
`mediainit /dev/rct`	*Initializes a cartridge tape for the specified device file*
`mediainit -i 2 /dev/r98286A`	*Initializes a flexible disk with interleave factor 2 for the specified device file*

6. If you get the following message:

   ```
   mediainit: this type of device unsupported
   ```

   you probably used the block special file instead of the character file.

7. When the initialization completes, you need to place a file system on the media unless you plan to use the media only for data storage or backups. See the later section named "Creating a File System".

Creating a LIF Volume

The LIF format is an HP standard format used for directories and files on flexible and hard disks. It is used on Series 200, 300, 500, and 800 computers that use the HP BASIC or Pascal operating systems, especially systems using 5.25- and 3.5-inch flexible disks made by Hewlett-Packard Company.

Prerequisites and Conditions

- If your media has never been initialized, you must use **mediainit** to initialize the media, then use `lifinit` to create a LIF volume.
- A LIF volume can be created directly on a disk or within the HP-UX file system.
- Test previously initialized media to see if it is a LIF volume *before* you initialize a disk. Assuming the media to check is associated with `/dev/rdsk/4s0`, execute:

  ```
  lifls /dev/rdsk/4s0
  ```

- You get one of the following responses:
  - For an initialized disk that is a LIF volume and contains files, you see a listing. For example, you might see:

    ```
    PLOT_DEMO HELLO_TEST MY_FILE
    ```

  - For an initialized disk that is a LIF volume and has no files, you see an empty line.
  - For an uninitialized disk, you see:

    ```
    lifls: Can't list /dev/rdsk/4s0; not a LIF volume
    ```

  You should continue, or not, depending on what you see.
- There is No SAM Procedure.

Manual Procedure for Making a LIF Volume

1. Become the root user.
2. Use `lifinit` to create a LIF volume according to the following syntax:

   `lifinit` [`-v`*nnn*] [`-d`*mmm*] [`-n`*VOL_NAME*] *FILE_NAME*

   If you use the last two parameters, capitalize them.

   a. The following example writes a LIF volume header to the disk associated with the device file named `/dev/rdsk/3s0`:

   ```
   lifinit /dev/rdsk/3s0
   ```

   b. The following example writes a LIF volume header (named WORK) to an HP-UX file (named TMP) that will be copied to a flexible disk where the volume size (270336) is the number of bytes contained on a 5.25-inch flexible disk and the disk can contain 240 directories.

   ```
   lifinit -v270336 -d240 -nWORK TMP
   ```

3. The *HP-UX Reference* manual has more information about the options used with `lifinit`.

Creating a File System

Except for streamed data, HP-UX cannot use media until you place a file system on it. The following cases suggest when you might create a file system:

- To move a directory such as `/users` onto another disk when your root file system no longer fits on the original disk.
- To facilitate backing up and restoring files on an optical disk.

Prerequisites and Conditions

- If you use SAM, the procedure is part of adding a hard disk drive, and you need not create a file system or initialize any media. If you work manually, create block and character device files for the drive.
- Pay careful attention to the type of device file when you use commands. For example, `newfs` requires the character device file while `mount` requires the block device file.
- The media must not currently have a mounted file system. Use `/usr/bin/bdf` to check this (the file system will not appear on your display).
- To use `/etc/newfs`, the drive having your media must have an entry in `/etc/disktab`. If the entry does not exist, you must make one. Use the directions given in `disktab`.
- The `/etc/disktab` file has many of the parameters for `newfs`. Examine this file before making the file system.
- The `newfs` command lists the alternate superblock locations in `/etc/sbtab`. Copy these locations for later use with `fsck` if your file system becomes corrupted.

- The `newfs` command automatically creates a file system having short file names (14-character limit). You can create a file system having long file names (up to 255 characters). If you create a "short" system, you can change it to a "long" system. You cannot easily change a "long" system to a "short" system. If you use SAM, you choose between having short and long file names.

SAM Procedure for Creating a File System

Creating a new file system is part of adding a hard disk drive.

```
System Administration Manager
    ↓
File Systems ->
    ↓
Add a Hard Disk Drive ...
```

At this point, enter data related to creating the file system. If you need help, call the Help Screen.

Manual Procedure for Creating a File System

If you need it, the following subsection has information about the syntax and parameters for newfs.

1. Log in as the root user; turn on the drive if necessary; and insert the media if necessary.
2. To make a file system using default values, go to the next step. If you need to change the default swap space, minimum amount of free space, or other parameters to `newfs`, continue as follows:
   a. Examine the information about parameters in the following section.
   b. Determine alternate values you need to use.
   c. Write the values down, if necessary.
   d. Go to Step 3 where you will specify them.
3. Use `/etc/newfs` to make the file system.
   a. The following example uses default values to make a file system on an HP7937 disk drive associated with the device file named `/dev/rdsk/1s0`:

      ```
      newfs /dev/rdsk/1s0 hp7937_noswap
      ```

   b. The following example creates the same file system with a non-default amount of swap space:

      ```
      newfs -s 48 Mbytes /dev/rdsk/1s0 hp7937
      ```

   c. The following example creates a file system on an optical disk in a Model 650A Optical Disk drive using a SCSI adapter.

      ```
      newfs /dev/rscsi hpS6300.650A
      ```

4. When the command finishes, mount the file system before attempting to use it. See the later section named "Mounting or Unmounting a File System".

Supplemental Information for Making a File System

Use `newfs` or `mkfs` to make a file system. HP recommends using the first command, and that procedure is documented here. The syntax is:

`newfs` [`-L` | `-S`] [`-n`] [`-v`] [*mkfs-options*] *device_file disk_type*

where the options work as follows:

-L	Creates a file system with long file names (up to 255 characters). The default is the standard 14-character file name limit.
-S	Creates a file system with short file names.
-n	Prevents the bootstrap programs from being installed.
-v	Provides verbose mode, which lists `newfs` actions.
mkfs-options	The options to `mkfs` that override default parameters. The `/etc/disktab` file defines many of these options. To get more information, see the *newfs*(1M) or *mkfs*(1M) entries in the *HP-UX Reference* manual.
device_file	The character device file name for the disk drive on which you are creating the file system.
disk_type	The type of disk as specified in `/etc/disktab`. The `newfs` command uses defaults from `mkfs` and `/etc/disktab`. The `mkfs` defaults appear in the table on the next page (N/A means Not Applicable). The `disktab` file contains disk-specific information for `newfs`. See the *disktab*(4) and *newfs*(1M) entries in the *HP-UX Reference* manual. If your disk is not listed in `/etc/disktab`, read the comments in that file to get information about creating a new entry.

The supplemental information continues on the next page.

Parameter	/etc/disktab Names	Range	Default	Comments
size	s0	N/A	none	If using `newfs` size is taken from `/etc/disktab` by default. Size is total disk size minus swap size in 1024 blocks.
block-size	b0	4K or 8K	MAXBSIZE (8K)	Specified in bytes.
frag-size	f0	1024 to block-size	1024 bytes	Specified in bytes (an even multiple of 1024).
number of tracks per cylinder	nt	greater than 0	16	Taken from `/etc/disktab` if using `newfs`.
number of cylinders per disk	nc	1 to 32	16	
% free space reserved	N/A	0 to 100	10 %	On crossing threshold, only superuser can continue writing.
revolutions per minute	rpm	N/A	3600	If using `newfs` parameter is revolutions/minute and value can be taken from `/etc/disktab`.
number of bytes per inode	N/A	1 to (function of file system size and other parameters)	max (2048, fragment size)	Number of allocated inodes is function of block size (2048 max inodes per cylinder group).

Information for making a file system continues on the next page.

For most installations, the default file system setup is correct. If you have an installation that requires a different file system configuration, use the following guidelines:

- If your system has many small files, you can decrease the average number of bytes per inode. This gives you more inodes, and lets you create more (but smaller) files. Having many inodes takes more space on your file system.
- If your system has a few large files, you can increase the space available for data by increasing the average number of bytes per inode.
- Decreasing the value of *minfree* lets you write to an additional percentage of file system space. The lower the percentage, the greater the possibility that your file's blocks will be scattered on the disk. Performance decreases as the disk fills up.
- Decreasing the file system size gives you more swap size, which lets you run larger programs, but decreases the area in which you can store files. Some optional application programs for HP-UX require above-average amounts of swap space. The applications' manuals should discuss swap space requirements.
- Increasing the file system size gives you less swap size. If you have a swapping device separate from your file system, this gives you a larger file system and still lets you execute large programs.

When you know how you want to create your file system, go back to the procedure and make the system.

Mounting or Unmounting a File System

HP-UX cannot use a file system unless it is mounted. This operation establishes a link between the file system and the device file for the mass-storage device. The procedure for mounting or unmounting a file system is similar for a hard disk, optical disk, CD-ROM, and flexible disk.

Prerequisites and Conditions

- You cannot mount devices attached to a client in an HP-UX cluster or execute commands related to mounting or unmounting a file system from a client. NFS mounts can occur from a client.
- Make sure the drive is correctly addressed and has an appropriate device file. If you use SAM, this is done for you.
- You should have initialized the media according to recommendations for the type of media. If you use SAM, this is done for you.
- You should have created a file system on the media according to the type of file system you need and the values for parameters that make the file system function appropriately. If you use SAM, this is done for you.
- Do not unmount a file system that has any open files. Use `ps -ef`, `fuser`, and `cfuser` to check this. Kill unnecessary processes. In an HP-UX cluster, you may need to save a person's work and log the person out.

SAM Procedure for Mounting a File System

Under `File Systems ->`, You can do this as part of adding a hard disk drive or adding an NFS file system. Also, you can work as follows:

`System Administration Manager`
↓
`File Systems ->`
↓
`Change File System Parameters ...`

At this point, opt to mount a file system now. If you need help, call the Help Screen.

SAM Procedure for Unmounting a File System

Under `File Systems ->`, you can do this as part of removing a hard disk or removing an NFS file system. Also, you can work as above for mounting a file system except you opt to unmount the file system by entering `n` instead of `y` for mounting the file system. Also, changing a hard disk drive address will unmount a hard disk.

Manual Procedure for Mounting a File System

`mount /dev/`*block_dev_filename mount_dir*

For example:

```
mount /dev/dsk/1s0 /dir1
```

Manual Procedure for Unmounting a File System

Close all open files on the device being unmounted.

`/etc/umount` *block_dev_filename*

For example:

```
/etc/umount /dev/dsk/1s0
```

Moving a File System onto Another Disk

Besides mounting or unmounting a file system, you may need to move a file system onto a new disk so as to create more disk space. The manual procedure assumes you mount `/users` (the file system) on a new disk you added to the system to create additional file space. The temporary new directory is called `dir1` (the new disk); and your device file is named `/dev/dsk/1s0`. You can extrapolate from the examples for different situations. There is no SAM procedure.

Manual Procedure for Moving a File System to a New Disk

1. You can be in run-level 2, the multi-user state, but it is better to do this work when no one is using the system.
2. Make a new, dummy directory for mounting the file system as follows.

   `mkdir` *mount_dir*

   For example:

   ```
   mkdir /dir1
   ```

3. Mount the new file system onto *mount_dir* as follows:

   `mount /dev/`*block_dev_filename mount_dir*

   For example:

   ```
   mount /dev/dsk/1s0 /dir1
   ```

4. Change to the directory where you will eventually mount `/dev/dsk/1s0` by executing:

   `cd /users` *(In the example for this procedure)*

5. Copy all the files from this directory to the new disk by executing:

   ```
   find . -hidden -depth -print | cpio -pvdumx /dir1 | tee /tmp/cpio.log
   ```

   You see the list of files and can see the log for the copy in `/tmp/cpio.log`. The copy takes an hour or so.

6. If desired, you can verify the copy, which takes awhile, by executing:

   ```
   dircmp /users /dir1 | tee /tmp/dircmp.log
   ```

7. Edit `/etc/checklist`, adding the following line:

   ```
   /dev/dsk/1s0 /users hfs rw 2 0 #/users mounted
   ```

8. Unmount the new disk by executing:

   ```
   umount /dev/dsk/1s0
   ```

9. You can shut down the system so you have no file activity, or you can use `/etc/fuser` to see if there is any activity. When there is no activity, remove everything in `/users` by executing the following two commands:

   `shutdown 30` *Wait until you get to single-user state*
   `rm -r /users/*` *Removes all directories and files in /users*

10. If necessary, reboot the system by executing `reboot` with no options. This returns you to multi-user state. Since you added an entry to `checklist`, the `/users` file system will be mounted on `/dev/dsk/1s0`, which was what you wanted to do. Now, you can remove `/dir1`, which finishes the cleanup.

Automatically Mounting/Unmounting File Systems

Your system created one root file system when you installed HP-UX. You can use SAM or a manual method to mount additional file systems. Another alternative is to have an entry in `/etc/checklist` that will cause `/etc/rc` to automatically mount additional file systems (one per disk) and provide secondary disks for swap space during the startup process. If you do this, the shutdown process uses `umount -a` to automatically unmount devices mounted by `checklist`. By doing this, you need not remember to mount the file systems you want.

Prerequisites and Conditions

- The file named `/etc/bccheckrc` checks each file system in `/etc/checklist` during start up and uses the clean byte to determine if the system was properly shut down. If a file system is corrupt, `bccheckrc` runs `fsck` during the startup process.
- When the file systems are clean, `/etc/rc` mounts the ones listed in `/etc/checklist`.
- Add a line in `/etc/checklist` for each file system you want to mount, check, or use for swapping. Do not have any blank lines.

SAM Procedure Mounting With /etc/checklist

System Administration Manager
↓
File Systems ->
↓
Change File System Parameters ...

At this point, you see the data-entry screen. Opt to mount on boot, which adds an entry to `/etc/checklist`. If you need help, call the appropriate Help Screens. You can also do this when adding a disk drive or an NFS file system.

Manual Procedure for Editing a Line in /etc/checklist

If you need an explanation of the parameters, the section following this procedure describes them.

1. Do not be in the single-user state prior to editing.
2. Edit lines in `/etc/checklist` as required, using the the following format:

   special_filename directory type options bckup_freq pass_number comment

   A file that mounts file systems, provides swap space, and mounts file systems on a networked system might look like this:

   ```
   # Sample /etc/checklist file. (see CHECKLIST(4)).
   /dev/dsk/0s0 /  hfs defaults 0 1 # 7937 96Mswap
   /dev/dsk/1s0 /users  hfs defaults 0 2 # 7937 430Musers
   /dev/dsk/1s0 /users  swap defaults 0 0 # 7937 96Mswap
   ww:/usr/lib/tex /usr/lib/tex nfs defaults 0 0 # nfs-mount
   ww:/users/ftp /users/ftp nfs defaults 0 0 # nfs-mount
   ```

3. When you finish editing, save the file. Later, during startup, HP-UX mounts the specified file systems because `/etc/rc` runs `mount -a`. Shutting down the system will unmount them because `shutdown` runs `umount -a`.

Specifying Parameters for /etc/checklist

The next several sections describe parameters used in each line in `/etc/checklist`.

First Parameter: special_filename

The block device name used by `fsck`, `mount`, or `swapon`. For example, `/dev/dsk/4s0` for a file system mounted on a disk having bus address 4.

Second Parameter: directory

The directory on which you mount the file system must already exist and you must specify an absolute path name.

Third Parameter: type

A code according to the following types of devices:

hfs — Specifies an HFS file system, the standard file system on HP-UX.

nfs — Use when *special_filename* is a remote NFS file system; use the serving machine name followed by a colon (:) and the path on the serving machine of the directory to be served. See the example used earlier in the procedure. NFS entries ignore the *pass_number* and *bckup_frequency* fields.

swap — Use when the *special_filename* will be used as a swapping device. During system startup, the `/etc/rc` file executes `swapon -a`, which enables devices labeled with type `swap`. The *directory, pass_number,* and *bckup_frequency* fields are ignored for lines having `swap`. The `mount` and `fsck` commands ignore lines containing `swap`. Do not have a swap entry for the root file system as it is assumed.

swapfs Use for dynamic swap.

ignore Use when a line should not be used by any command, but you want to keep the line.

Fourth Parameter: options

If an entry is marked as `hfs` or `swapfs` and is used by the `mount` command, this field can contain a list of comma-separated options to the `mount` command.

Fifth Parameter: bckup_frequency

This field is reserved for future enhancements to backup utilities.

Sixth Parameter: pass_number

The `fsck` command , when used with `-P`, uses this field to determine the order for checking file systems.

The root file system (`/dev/dsk/0s0`) should be 1 and should be the only file system set to "1". A file system set to 2 is checked after the root file system; a file system labeled 3 is checked after 2; and so on. When more than one file system has the same value, `fsck` deals with all of them in parallel, which shortens the time required for the check. The command ignores file systems set to 0. If you do not set a value, the command checks each file system sequentially.

Seventh Parameter: comment

Precede this field by a # and add your comment.

Checking File System Consistency

The HP-UX file system can develop inconsistencies over a period of time. Turning off the computer without previously shutting down the file system corrupts the system. Except for obvious events (having a power failure), you cannot relate file system corruption to specific problems.

Check the file system for consistency periodically and anytime you suspect a problem. Some commands such as `/etc/update` will not function properly unless you have a clean file system.

Prerequisites and Conditions

- The `fsck` command has options you may want to use. See the *fsck*(1m) entry in the *HP-UX Reference* manual if you need information about how file system consistency checks work.
- The `fsck` command can display messages and prompt you to answer questions. To respond properly, you need to understand the file system. See the *HP-UX System Administration Concepts* for conceptual information.
- In an HP-UX cluster, do *NOT* run `fsck` from a client.
- If you had an unclean shutdown, starting up HP-UX runs `fsck` in a preen mode. The `bcheckrc` entry in `/etc/inittab` causes this.
- Your system was shut down improperly if you did not use the `shutdown` or `reboot` commands correctly prior to turning off the computer.
- The root file system must contain `/lost+found`. The `fsck` command uses this directory. Look for "lost" files in the directory when the check finishes.

- You must run `fsck` in an appropriate run-level and use an appropriate device file:
  - To check the root file system, which remains *mounted,* run `fsck` on the block device file (probably `/dev/dsk/0s0`).
  - To check a file system on another device, which must be *unmounted,* run `fsck` on the character (raw) device file (possibly `/dev/rdsk/1s0`).
  - Use the single-user run level. Do not run a file system check in a multi-user run level.
- You should probably not run `fsck` from `cron` because the check is meaningful only when the system is quiet during a check.
- Unless you specify a device file, `fsck` checks all `hfs` entries in `/etc/checklist`.
- If you suspect corruption, you should probably not run `fsck` with the `-y` option. You may not want to ever do this because it is usually better to see the messages and respond interactively.
- If you have a corrupted system and `fsck` does not fix it, you may need to rebuild the system by using your recovery program or reinstalling HP-UX from your original media. You can usually prevent having to do this by checking the system periodically and working carefully when managing the file system in the single-user state.

Manual Method for Checking Consistency

1. To prevent possible problems, be in the single-user state (run `/etc/shutdown`). You must be in the single-user state to check the root file system. Other file systems need only to be unmounted.
2. If necessary, change to the root directory (`cd /`).
3. Use `ps -ef` and examine the listing. Check that the unmounting process is not still running. You may need to repeat this a few times because the unmounting process can take a few moments. You know everything is unmounted when the process no longer appears in the listing.
4. Run the `fsck` command. The following examples show possibilities:

`fsck`	*Checks the root system, or checks systems having hfs entries in /etc/checklist*
`fsck /dev/dsk/0s0`	*Checks the root file system in an interactive mode. This lets you answer questions about fixing problems.*
`fsck -p /dev/dsk/0s0`	*Checks the root file system in preen mode, which can repair some corruption and does not remove any data*
`fsck -y /dev/rdsk/1s0`	*Checks a file system on a non-root disk and attempts to fix all problems*

5. The command moves through seven phases (passes 0-6). As the command runs, you may get messages related to the following items:

When You Get This ...	**Consider Responding Like This ...**
Unreferenced inodes	Continue the check, noting the unreferenced nodes for future reference.
Unreferenced files and fifos	Continue the check, noting the items for future reference.
Link counts in inodes too large	Elect to fix them.
Missing blocks in the free list	Elect to fix them.
Blocks in the free list also in files	Elect to fix the situation.
Wrong counts in the superblock	Elect to fix them.
Clean byte marked wrong	Fix the byte.

   Respond to particular situations according to your best view of how to continue during the check, and perhaps fix problems when the check finishes.

6. When the check finishes, the system should start up again, indicating you have a clean file system. If the system does not automatically start up, execute `reboot`. You may get a message indicating how to execute `reboot`. In any event, not starting the system up immediately can negate the work done by `fsck`.

7. If a file system consistency check cannot repair your system, you have some alternatives for getting the system going again:

   a. Use your recovery system to boot a partial system. Then, restore the system by getting files from your backed up file systems.

   b. Install the system again from your original media. Then, restore your customizations and users from backed up file systems.

Adding or Deleting Swap Space

When you installed HP-UX, you set up a certain amount of swap space. Later, you may need to increase (or decrease) this amount. The task is not trivial. Manually, it consists of completing many tasks.

Prerequisites and Conditions

- Swap space is an integral part of your HP-UX kernel. If you do not understand swap space, see the *HP-UX System Administration Concepts* manual.
- If you have only one disk and you must increase the swap space, you can use your original tape to reinstall HP-UX and specify an amount of swap space that is greater than the default amount, or you can use dynamic swap space.
- You can get temporary relief by changing the mix of running programs to reduce virtual memory activity. This will not solve your problem.
- You cannot free swap space until a process completes. This can be a problem on HP-UX clusters when clients spawn shells and processes all day without killing them.
- While your system runs, you can get the following message:

  `swap: rmap ovflo, lost`*blocknumber*`,`*blocknumber*

  Read the information following the message because, later, you need to reconfigure the kernel with the next higher values for `dmmax`, `dmtext`, and `dmshm`.
- Before you install an additional disk for swap space, set up an appropriate device file and make sure the kernel has a device driver for the disk.
- Determine the amount of swap space you need before you begin this process. You can use `/usr/bin/sam/bin/swapinfo` to determine how much space you have.
- You can get temporary relief by enabling dynamic swap space.

SAM Procedure for Adding Swap Space

`System Administration Manager`

↓

`Kernel Configuration ->`

↓

`Add to the Swap Configuration ...`

At this point, you see the date-entry screen. Enter appropriate values according to the fields. If you need help with a particular field, call the Help Screen for that field.

For a client in an HP-UX cluster, you can add swap space while adding a disk drive. This action also adds swap to the kernel.

SAM Procedure for Deleting Swap Space

`System Administration Manager`

↓

`Kernel Configuration ->`

↓

`Delete from the Swap Configuration ...`

At this point, you see the date-entry screen. Select the swap entry you want to remove and press `Perform Task`.

Removing a hard disk removes a swap entry from `/etc/checklist`, but not from the kernel. You may want to do both.

Manual Method for Adding or Deleting Swap Space

Some steps may send you to other parts of this manual to complete a task. If you do this, complete the task and return to these steps. Some tasks can require being in the single-user state.

1. Initialize the media on the disk. See the section named "Initializing HP-UX Media" if you need help with this.

2. If you plan to use the disk entirely for swap space, skip this step and go to the next step. Otherwise, using the amount of swap space you determined, create a file system on your disk, knowing that the amount of swap space set up on the disk must equal or exceed the amount you calculated. If you use the default amount for the new disk, this value should equal or exceed your calculated amount. See the section named "Creating a File System" if you need help with this.

3. Edit your `/etc/conf/dfile`, adding the appropriate entries. This file (the configuration description file) will have a section like the following:

   ```
   * SWAP CONFIGURATION
   swap devname address swplo [nswap]
   ```

   If you add swap space, you must have a line for the root disk that might look like this:

   ```
   swap cs80 0e0000 -1  #rootdisk
   ```

   Under that line, add the line for the second swap disk. Assuming you added a similar disk at bus address 2, add a line such as the following:

   ```
   swap cs80 0e0200 0  #swapdisk
   ```

   If you need help with this, see information about the dfile in the chapter named "Reconfiguring the Kernel".

4. Edit `/etc/checklist` if you want to automatically add the disk as swap space during startup. See the earlier section named "Automatically Mounting/Unmounting File Systems". If you do not want to do this, skip this task. If you skip it, you must manually execute **swapon** *swap_dev_filename* after you reconfigure the kernel and every time you restart the system to enable the swap space.
5. Run `/etc/config` to rebuild the kernel. See the chapter named "Reconfiguring the Kernel" if you need help with this.
6. If you did not edit `/etc/checklist`, which automatically sets up swapping to the second disk, use the **swapon** command to set up swapping. See the *swapon*(1M) entry in the *HP-UX Reference* manual to get more information.
7. Removing swap space amounts to removing the lines you included in `/etc/conf/dfile` to add the swap space, rebuilding the kernel, and rebooting the system.

Adding Dynamic Swap to a File System

Dynamic swap allows processes to use space in the file system as overflow swap space when their need for swap space overflows the swap area. This swap space exists in addition to the configured primary and alternate swap space and does not replace them. The system uses available file system blocks for swap space. The system uses existing fixed swap space before it uses the dynamic swap space.

Prerequisites and Conditions

- Unlike changing fixed swap space, you can enable dynamic swap without creating a new kernel.
- Enabling dynamic swap allows more processes to run simultaneously. This can slow down the system response time.
- Use `/usr/bin/bdf` to check the available space on the file system.

  You could use the `View Disk Space` screen provided by SAM under `File Systems ->`. You could see something like this:

  ```
  Filesystem          kbytes    used   avail capacity Mounted on
  /dev/dsk/0s0        439998  234984  161014    59%   /
  /dev/dsk/2s0        390975  127157  224720    36%   /other
  /dev/dsk/1s0        439998  240286  155712    61%   /users
  ww:/usr/contrib/hptag
                      487022  388642   49677    89%   /usr/contrib/hptag
  ww:/users/ftp       487022  388642   49677    89%   /users/ftp
  ```

  Do not enable dynamic swap on file systems that are heavily used, such as root (/) or `/users`.

- File system swap performance is slower than device swap. If a system performs noticeably slower, you should provide an additional swap device.
- You can enable dynamic swap only on local file systems.
- The *nswapfs* parameter determines the maximum number of file systems to which you can swap.
- Do not enable dynamic swap space on more than one directory per file system. The swapper will compete for swap space out of the same file system. This will be very inefficient.
- The *HP-UX System Administration Concepts Manual* has information about the limitations of using dynamic swap.

SAM Procedure for Enabling Dynamic Swap

`System Administration Manager`

↓

`File Systems ->`

↓

`Add Dynamic Swap ...`

At this point, you see the data-entry screen. Enter appropriate data in each field. If you need help with with a particular field, call the Help Screen for that field.

Manual Procedure for Enabling Dynamic Swap

Before you enable the dynamic swap space, you may want to examine *swapon*(1M) in the *HP-UX Reference* manual. In addition, if you need information, the following section explains the syntax and parameters. When you feel comfortable with the `swapon` command, use the following procedure to enable the space:

1. You must be the root user, but you need not be in the single-user state.
2. To enable dynamic swap on `/usr` with:

   a. a minimum of 4096 file system size blocks reserved for swap
   b. a maximum of 8192 file system size blocks allowed for swap
   c. 1024 blocks reserved for file system use
   d. swapping priority of 1

   you would execute:

   ```
   swapon /usr 4096 8192 1024 1
   ```

You can examine `/etc/disktab` to determine the file system block size.

You should probably consider the following additional steps:

1. If you want to enable dynamic swap automatically each time you start up the system, include the following entry in `/etc/checklist`:

   ```
   default /swap swapfs min=10,lim=4500,res=100,pri=0 0 0
   ```

2. When you have enabled dynamic swap, you can check the file system for the amount of file system space being used for swapping by executing the following command:

   ```
   bdf -b /swap
   ```

   The output could look like this:

   ```
   Filesystem          Kbytes    used     avail    capacity  Mounted on
   /dev/dsk/1s0        121656    16464    93024      18%       /swap
   Swapping            120832    512      38400      19%       /swap
   ```

Information About the Swapon Syntax and Parameters

swapon *mount_point* or *special_file* [*min*] [*max*] [*reserve*] [*priority*]

The parameters work as follows:

mount-point
: The directory in which the file system enabled for dynamic swap is mounted (for example, /usr).

special_file
: The special (character) device file for the partition that is enabled by dynamic swap (for example, /dev/rdsk/1s0).

min
: The number of blocks reserved for the swap system to use. The minimum number of blocks reserved is determined by the value of the dmmax parameter, which is described in the appendix on operating system parameters. The default minimum number of blocks reserved is 512 (.5 Mbytes).

max
: The number of blocks the swap system can take from the file system. The default for *max* is 0 (take as many blocks as needed).

reserve
: The number of blocks reserved for file system use. The default for *reserve* is 0 (no blocks reserved for the file system).

priority
: The order in which systems will be used for swapping; 0 (the default) is the highest interleaving priority and 10 is the lowest. The file system having the lowest priority is the first one used for swap space.

The *HP-UX System Administration Concepts Manual* has additional information about the swapon command.

Transferring Files

You will probably spend a lot of time transferring files among directories and systems. Unlike most modules that explain a single procedure, this module describes situations and relates them to methods for transferring files.

Prerequisites and Conditions

- Each command used to transfer files has particular features. If you have an HP-UX cluster, for example, using `find` with the `-depth` and `-hidden` options catches the CDFs. In another situation, you may need raw file I/O, which you could do with `cpio`, `dd`, `ftio`, or `tar`. If you need a certain capability and you do not understand a command, take time to study its entry in the *HP-UX Reference* manual.
- Use commands such as `cp` and `mv` to copy and move files among directories on your system. Beyond using `mkdir` to create necessary directories, you have no significant prerequisites.
- Transferring files to remote systems via a network service or a communication utility such as UUCP requires special hardware and some customization of your system. In many situations, you need to cooperate with the system administrators of remote systems.
- The procedures described in a later chapter named "Backing Up and Restoring the System" explain the commands used to transfer files from an HP-UX file system to mass-storage media such as cartridge tapes and optical disks. Read that chapter to learn how to use commands such as `fbackup`, `backup`, `cpio`, `tcio`, `find`, `ftio`, `dd`, and `tar`.

- This module focuses on transferring files among systems. In this context, you have the local system and another system administrator has a remote system. Two systems can be connected via:
  - A LAN cable and some networking service (for example, ARPA).
  - A direct RS-232-C connection (UUCP).
  - An indirect modem/phone connection (UUCP).

  A remote system can be another HP-UX system, a UNIX ™ system, a UNIX-like system such as HP/Apollo, or an MS-DOS ™ system. You may encounter other remote systems.
- Given all the possibilities, you need to assess your needs. Then, it usually works best to choose certain methods and learn to use them well.

Transferring Files via Networking

Using a network service lets you transfer files among systems connected by a Local Area Network (LAN). HP-UX has many LAN services (for example, NS-ARPA, NFS, RFA). Commands in one service may duplicate the apparent functionality of commands from another service. To see possibilities, read the documentation for each service and examine the detailed information about the features of its commands.

Prerequisites and Conditions

If you have not set up networking on your system, see the module named "Networking Systems" in the chapter named "Constructing an HP-UX System". If you have not yet done anything, this can take some time and effort. SAM can do many of the tasks.

After you set up a networking service, you can use it to transfer files among HP-UX and UNIX ™ systems. Each service provides certain commands. Therefore, you need to choose among services and commands.

NS-ARPA Services
ARPA/Berkley Services
ARPA Services/Vectra PC

You get the following file-transfer services:

- File Transfer Protocol (ftp)
- Telnet (telnet)
- Simple Mail Transfer Protocol (SMTP)

The services provide commands for file transfer and mail (for example, `ftp`, `rcp`, and `rlogin`. Additional commands include `remsh`, which provides a remote shell in which you can transfer files using conventional commands. The `sendmail` facility provides commands such as `mailx`, which can send electronic mail and transfer files as well.

See the *Using ARPA Services Manual* and the *Installing and Maintaining NS-ARPA Services* manuals to get complete information.

Network File System (NFS)

You get the following file-transfer services:

- NFS Remote File Access
- Remote Procedure Calls (RPC)
- External Data Representation (XDR)
- Yellow Pages (YP)
- Virtual Home Environment (VHE)

Commands for file transfer relate to NFS. See the *Using and Administering NFS Services* manual and the *Programming and Protocols for NSF Services* manual to get complete information.

HP AdvanceNet on HP 9000 LAN

You get the following file-transfer services:

- Remote File Access. RFA connects your file system to a remote file system via network special files in a directory named `/net`.
- Network File Transfer. NFT provides commands such as `dscopy`. The file type is a stream of bytes.
- Link Level Access.

HP AdvanceNet also provides services related to Series 500 computers. See the *LAN User's Guide* and the *Using Network Services* manual to get complete information.

Transferring Files via UUCP

The UUCP facility has existed on UNIX ™ and HP-UX for a long time. It is a traditional method for transferring files among systems that use RS-232-C protocols.

If you have a RS-232-C direct connection or a modem connection with another HP-UX or UNIX ™ system, you can use the UUCP facility to transfer files. If you do not have these connections, you need to set them up. This takes some time for planning, installation, and testing. The *UUCP Concepts and Tutorials* manual describes what you need to know about installing, customizing, and using `uucp`, and explains its related commands (`cu` and `uucico`). SAM can help you perform many of the tasks.

Transferring Files to Systems Using the LIF Format

The LIF utilities provide commands for reading and writing files in the HP Logical Interchange Format (LIF). The HP BASIC and HP Pascal workstations utilize LIF.

Getting a set of HP-UX files into a format a language workstation can read requires three HP-UX commands: `lifinit`, `lifcp`, and `cat`.

1. Create a LIF volume with `lifinit`.
2. Use `lifcp` to write the files to this volume.
3. Use `cat` to write the LIF volume to the initialized LIF media (often a flexible disk).

Copying HP-UX Files to LIF Volumes

Once you create a LIF volume, you can write files to the volume using the `lifcp` command.

`lifcp` *hpux_file VOL_NAME*:*FILE_NAME*

The parameters work as follows:

hpux_file	A file in your HP-UX directory.
VOL_NAME	The name you gave to the LIF volume when you created it.
FILE_NAME	The name that file is given on the LIF volume.

Use uppercase letters for LIF names.

For example, to transfer the HP-UX file called `testing.p` (located in `/users/engel/progs`) to the LIF volume called `VOL1` in the current directory, execute:

```
lifcp /users/engel/progs/testing.p VOL1:TESTING
```

You get an error message when the LIF volume has insufficient room for the file. You must then create another LIF volume and copy the file to it.

Moving the LIF Volume to Flexible Disk

After you copy files to the LIF volume file, use the HP-UX `cat` command to write the LIF volume to the flexible disk.

1. Insert the flexible disk into the disk drive.
2. List the contents of the disk by executing:

   `lifls /dev/`*file_name*

   where *file_name* names the raw device file associated with the disk drive holding your flexible disk (possibly, `/dev/rfd9127`).

   Listing the contents of the disk helps you identify files you may want to save later.
3. Use `cat` to copy the LIF volume file to the disk. This action overwrites everything on the disk.

   If your LIF volume is named `VOL1` on HP-UX and on the flexible disk, and if the device file for the disk is `/dev/rfd9127`, execute:

   ```
   cat VOL1 > /dev/rfd9127
   ```
4. When the copy finishes, remove the LIF volume from your current directory by executing:

   ```
   rm VOL1
   ```

Adding Files to a LIF Volume

Assume you have a 5.25-inch flexible disk in LIF ASCII format and you want to write an additional file to the disk, leaving the current contents of the disk intact. Execute:

`lifcp` *hpux_file* `/dev/`*file_name*`:`*FILE_NAME*

hpux_file — The HP-UX file you want to copy to the LIF disk.

file_name — The name of the device file associated with the flexible disk drive (for example, `rfd9127`).

FILE_NAME — The uppercase file name given to the *hpux_file* stored on your LIF disk.

Before you copy a file to the disk, determine the storage space on the disk by executing:

`lifls -l /dev/`*file_name*

where *file_name* names the device file associated with the flexible disk drive. This lists the files on the disk and the space they use.

Transferring LIF Files onto HP-UX

You could need to copy LIF files from a flexible disk to the HP-UX file system (for example getting the latest revision of a program on a flexible disk and needing to transfer it to the hard disk on your system).

1. Place the flexible disk in the disk drive.
2. List the files on the flexible disk using the `lifls` command, as previously explained, to verify the file is on the disk.
3. Copy the file *FILE_NAME* from the disk into the HP-UX file *hpux_file* by executing:

   `lifcp /dev/`*dev_file:FILE_NAME hpux_file*

   where *dev_file* is the special device file name associated with the disk drive holding the LIF file.

   For example, to copy the file named `TESTING` from the disk in the disk drive associated with the `/dev/rdf9127` file, to a file called `testing.p`, execute:

   ```
   lifcp /dev/rdf9127:TESTING testing.p
   ```

Transferring Files to Systems Using MS-DOS

You may need to transfer an MS-DOS ™ file to HP-UX, and conversely. To discuss this, this section makes some assumptions:

- You have an HP Vectra computer running MS-DOS.
- The Vectra is connected to the HP-UX system via LAN.
- You have installed the ARPA services for PCs.

Transferring Files from a Vectra to HP-UX

Assuming you work on the Vectra, proceed as follows:

1. Start up DOS.
2. If necessary startup networking on the Vectra.
3. Insert the flexible disk containing the files you want to transfer into Drive A:.
4. Use `telnet` to log into the HP-UX system (for example, a system named `hpfcma`):

`telnet hpfcma`	*Starts the connection to the HP-UX system*
`login:johnb`	*Login as a user on the system. You must have a login there.*
`password:xxxx`	*Supply your password.*

5. While you are remotely logged into the HP-UX system, make a directory for receiving the files. For example:

   ```
   mkdir /dos
   ```

6. Exit `telnet` (usually with Ctrl-D).

7. Use `ftp` for the file transfer. See Ftp in the *Using ARPA Services* manual if you do not know how to use this facility. As reminders while using `ftp`, note the following items:
   a. Change to the directory getting the files (for example, `cd /dos`).
   b. To transfer non-binary files, use Ftp's `put` command.
   c. Use the binary mode for any executable DOS files and then use the `mput` command for executable DOS files. Later, to run these programs on HP-UX, you would need to run an applications such as the SoftPC or DOS-Coprocessor.
8. Repeat the above procedure for any additional disks.

Transferring Files from HP-UX to a Vectra

Assuming you work from a Vectra, work as you would in the earlier procedure (transferring files from DOS to HP-UX) with the following exceptions:

1. You need not do any steps involving `telnet`. Instead, make any necessary directories for receiving the HP-UX files on your DOS system.
2. You could not run any executable HP-UX files on the DOS system, so do not transfer them to the DOS system.
3. To transfer files, use `ftp` as before except that you use `get` instead of `put`.

Transferring Files to Non-HP Systems

The ability to "connect" systems via LAN, RS-232-C, or HYPERchannel ™ means you can transfer files among non-HP systems. Transferring files among these systems is mostly a problem-solving process that cannot be reduced to a stepped procedure. In general, you can:

- Become familiar with HP-UX file transfer commands and processes.
- Study the services provided for the various systems.
- Study the documentation for the various systems.
- Visit with systems experts.

In learning how to transfer files to non-HP systems, you may encounter any of the following services, networks, or systems:

- X.25
- NS VAX ™
- DECnet ™
- SNA Gateways
- ARPANET (or Defense Data Net)
- HYPERchannel ™
- LAN

Automating Processes (Commands)

The HP-UX system has a clock daemon named `cron`. This daemon can regularly (automatically) execute commands placed in a file that you create with the `/etc/crontab` command.

Prerequisites and Conditions

- Plan the processes (commands) you want to automate. This usually includes backups of certain file systems, but you may also want to regularly execute a program such as `sendmail.trim`.
- You should understand `cron` and `crontabs` before you automate a process. See the *cron*(1M) and *crontab*(1) entries in the *HP-UX Reference* manual.
- You automate a process by entering a line having six fields in a crontab file.
  1. Minute from 0 to 59.
  2. Hour from 0 to 23.
  3. Day of month from 1 to 31.
  4. Month of year from 1 to 12.
  5. Day of week from 0 to 6 (0=Sunday).
  6. A string that specifies a shell and a command to execute in the shell.
- Specifying a line also has the following conventions:
  1. Type a space or tab to separate fields.
  2. Use an asterisk to allow the entire range of values for a field.
  3. Specify a range with an initial and final value, separated by a dash.
  4. To specify several values for a field, enter the values separated by commas.
- Use `crontab -l` (the letter, not the numeral) to list the contents of an existing crontab file. If one exists, the file is probably:

  ```
  /lib/spool/cron/crontabs/root
  ```

  If you run UUCP, you probably also have:

  ```
  /lib/spool/cron/crontabs/uucp
  ```

Procedure for Automating a Process

1. Become the root user.
2. If your crontab file exists (/usr/spool/cron/crontabs/root), you can see the currently automated processes by executing:

   ```
   crontab -l
   ```

   You can automate another process by editing the file. For example, the following line appends kernel diagnostic messages every to minutes to the file named messages.

   ```
   05,15,25,35,45,55 * * * * /etc/dmesg - >>/usr/adm/messages
   ```

   If your system has contributed files, the following line automatically runs the sendmail.trim program 55 minutes after 11 pm every day.

   ```
   55 23 * * * sh /usr/contrib/etc/sendmail.trim
   ```

3. If you do not have a crontab file, you can create one (/usr/spool/cron/crontabs/root) by executing:

   ```
   crontab croninfo
   ```

   Add lines to root as explained in the above step.
4. When you finish editing the file, executing the following command ensures that the processes will be activated:

   ```
   crontab /usr/spool/cron/crontabs/root
   ```

Monitoring the File System

While it is not really alive, the HP-UX file system can seem like a living entity. Some files grow without bound unless you edit them. The system can add files without your knowing it. Some commands create files in certain situations.

Instead of describing a procedure, this module describes the files and directories you can examine to monitor the file system.

Prerequisites and Conditions

- Depending on how you customized HP-UX, the system may not have some of the files described later in this module.
- Monitoring the system and controlling the system relate as follows:
  - Monitor the file system to see what is happening to it and to see what people are doing.
  - Based on the information you get from monitoring the system, you subsequently control disk usage via setting permissions, removing files, visiting with users, setting up accounting procedures, and using memos to communicate restrictions.
  - On an HP-UX cluster, `wtmp`, `btmp`, `sulog`, and `shutdown` are CDFs. Pruning these files requires pruning all the elements of the CDFs, not only the file accessible from a given system. In addition, the `/usr/adm` directory is a CDF, not only the files in the directory.
- Your system may have files not mentioned in the list if you run special applications. You should examine documentation for your applications to see if they create log files.

Commands for Getting System Information

The following items explain how to get information about disk usage. Some commands generate long lists, which you may want to redirect to a file or pipe to `more`.

- Use `du` often to spot users who are increasing their disk usage.

  `du -s /users/*` *Displays the users with information about usage.*

  `du -a / | sort -nr` *Lists directories and file, largest first.*

- Use `find` to locate large or inactive files.

  `find / -mtime +90 -atime +90 -print > aging-files` *Records files not written or accessed in 90 days to an aging file.*

  `find / -size +1000 -print > big-files` *Records files larger than .5 Mbytes to a big file.*

- For a modest fee, you can obtain `/usr/contrib/bin`. This directory contains utilities such as `monitor`, which is a program that lets you examine the state of the system in terms of users, memory, CPU activity, swap space, clients, and so on.

The information continues on the next page.

Determining Free Space On a Disk

Execute `bdf` or `df` to see the amount of free space left on a volume (disk).

`df -t` *Provides a detailed report of disk utilization.*

For example, you might see:

```
/users/ftp (ww:/users/ftp  ):  139592 blocks             -1 i-nodes
                               974044 total blocks       -1 total i-nodes
                               737046 used  blocks       -1 used i-nodes
                                   10 percent minfree


/users     (/dev/dsk/1s0   ):  288480 blocks         128796 i-nodes
                               879996 total blocks   147456 total i-nodes
                               503516 used  blocks    18660 used i-nodes
                                   10 percent minfree

/other     (/dev/dsk/2s0   ):  456200 blocks         123897 i-nodes
                               781950 total blocks   131072 total i-nodes
                               247554 used  blocks     7175 used i-nodes
                                   10 percent minfree

/          (/dev/dsk/0s0   ):  333174 blocks         129233 i-nodes
                               879996 total blocks   147456 total i-nodes
                               458822 used  blocks    18223 used i-nodes
                                   10 percent minfree
```

While the examples provide suggestions, HP-UX has many commands for exploring the system. In general, you can find them in Section 1M of the *HP-UX Reference* manual.

HP-UX Log Directories and Files

By Looking Here ...	You Get This Information ...
/tmp	A directory in which users (and some commands) place files. Treat them as *temporary* files.
/lost+found	A directory used by `fsck` to place detached files.
/usr/preserve	A directory containing editor files saved after a system crash.
/etc/newconfig	Receives customizable files and shell scripts during an update so the files on the system do not get overwritten.
/usr/adm/sulog	Logs history of `su` command use.
/etc/wtmp	A binary file containing a log history and date changes. Use `last` to read the file.
/etc/btmp	A binary file containing a failed login history. You must create the file to use it. Use `last` to read the file.
/usr/lib/cron/log	Contains a history of actions taken by `cron`.
/usr/adm/shutdownlog	Contains a log of shutdowns.
/etc/backuplog	Contains a log of backups. Your system may have some backup log files in the same directory that have slightly different names.
/usr/spool/lp/log	Contains a log of LP spooler requests.
/usr/spool/uucp	The directory contains several files having information about using `uucp`.

Enabling Long File Names

For compatibility with the AT&T System V UNIX ™, HP-UX releases before 6.2 limit file length to 14 characters. An installed system defaults to 14-character file names, but you can enable long file names on a per-file-system basis. Then, a file name can have up to 255 characters (a feature introduced with the BSD 4.2 UNIX ™).

Prerequisites and Conditions

- In a short filename system:
  - You can enter a long name, but the system truncates it to 14 characters.
  - Directory names cannot exceed 14 characters.
  - Directory entries align on 32-byte boundaries because an entry contains 32 bytes of information.
- In a long filename system:
  - Entering a name longer than 255 characters that passes to a system call causes an error.
  - Directory entries can vary in size, guaranteeing alignment only on 4-byte boundaries.
  - The magic number is in the superblock and differs from that of a short filename system.
- Do not enable long file names unless you know why you need them. Having short file names helps ensure compatibility with other systems. It is difficult to convert back to short file names, and you can lose files in doing so.

- Applications that assume directories having an array of fixed-size entries do not work. You may need to recode directory routines (those using `ndir` or `dir.h`). See *directory*(3c) in the *HP-UX Reference* manual.
- To run applications compiled on releases not having long file names, you may need to visit with the developers.
- HP-UX commands, utilities, tools, and supported applications run with either type of file name.
- Backup your existing system before converting a file system.

SAM Procedure for Enabling Long File Names

You must first log in as the root user and be in the single-user run level if you plan to convert the root file system.

`System Administration Manager`

↓

`File Systems ->`

↓

`Convert File System to Long File Names ...`

At this point, you see the data-entry screen. Enter appropriate data in each field. If you need help with a particular field, call the Help Screen for that field.

Manual Method of Enabling Long File Names

1. Get the system into the single-user state if you convert the root file system.
2. You have two choices for enabling long file names:
   a. To have the system prompt you about enabling long file names for file systems in `/etc/checklist`, execute:

      ```
      /etc/convertfs
      ```

      You see messages and a question about continuing. Then, the program asks about converting all file systems. You should convert all or none of them. Saying no causes the command to prompt about converting each file system. HP does not recommend having mixed long and short file names on the same system.

   b. To convert a specific file system, execute the command and specify the character special device file for the drive having the filesystem, for example:

`/etc/convertfs /dev/rdsk/2s0`	*Converts the file system associated with the specified device file, probably the disk at bus address 2.*

3. When the command and the subsequent file system check finishes, mount the file systems. The system may reboot HP-UX if you converted the root file system.

Some Followup Steps

After converting to long file names, consider the following items:

- If you converted the root file system, `convertfs` reboots the system. Changes made to the file system superblock will not be overwritten by an update of the superblock in the system memory.
- Recompile all pre-6.2 (Series 300) programs that use routines listed in *directory*(3c) in the *HP-UX Reference* manual.
- If you need to convert back to short file names, the process can be tedious, but you can proceed as follows:
  1. Examine all file names to make sure they are 14 characters or less. Use `mv` to move files having excessively long names to files having appropriate names.
  2. Backup your entire file system once the names have been shortened.
  3. Recreate the file system with short file names by executing `newfs` with the `-S` option.
  4. If the root file system must be reconverted, reinstall it from the installation tape.
  5. Recover your previously shortened files from the backup media.

Writing Programs In Systems Having Long File Names

If you convert to long file names, you may need to examine programs and shell scripts that do not function properly. This section contains some guidelines and examples for doing this work. You will need to understand the C programming language and know how to use it.

Make sure shell scripts and programs do not assume a 14-character limit. Note that scripts and programs making this assumption would work on a system with short file names but could stop working when moved to a system with long file names.

User programs may need to determine if a file system uses long file names. To do so:

1. Create two unique temporary file names. One has 14 characters (file one) and the other has 15 characters (file 2). The first 14 characters of these two files must be the same.
2. Stat each file.
3. Compare `st_dev` and `st_ino` of the stat structure for the two files. If both fields of both files are the same, the file system in which file 1 and file 2 reside uses short file names. If they are not the same, it is a long file name file system.

Do not assume that file systems accessed via networking have a maximum length of either 14 or 255. Although most networking servers support one of these two maximum file name lengths, some vendors use a different size.

The next page begins a series of examples designed to help you solve problems related to running programs on a file system using long file names. To work through the examples, you may need to see numerous other documents.

Examples That Check File Name Length

The following pieces of C code will check file name length:

```
 /* WARNING: Although most UNIX (Trade Mark of AT&T) systems support file
     systems whose maximum file name length is either 14 or 255
     characters, there are some vendors who support file systems with
     various file name lengths.  If an application is to run in a NFS
     environment, is_truncated() should be used to determine if a
     particular length is allowed in a certain file system.
     is_lfn() and is_truncated() are written without error checking,
     users might want to check return code of system calls.
*/
#include <sys/types.h>
#include <sys/stat.h>
is_lfn()
{
 struct stat stat1, stat2;
 char file1[MAXNAMLEN+1];
 char file2[MAXNAMLEN+1];
 int i;

 sprintf(file1, "abcdefgh.%-.5d", getpid());
 sprintf(file2, "abcdefgh.%-.5da", getpid());
 creat(file1, 600);
 creat(file2, 600);
 stat(file1, &stat1);
 stat(file2, &stat2);
 if ((stat1.st_dev == stat2.st_dev) && (stat1.st_ino == stat2.st_ino))
  i = 0;
 else  i = 1;
 unlink (file1);
 unlink (file2);
 return(i);
}
```

The code continues on the next page.

```
/* This routine tells caller if the file can be created with that name (without
   being truncated). */

is_truncated(filename)
register char *filename;
{
 struct stat stat1, stat2;
 char other[MAXNAMLEN];
 int len, i;

 len = strlen(filename) - 1;
 strncpy(other, filename, len);

 creat(filename, 600);
 creat(other, 600);
 stat(filename, &stat1);
 stat(other, &stat2);
 if ((stat1.st_dev == stat2.st_dev) && (stat1.st_ino == stat2.st_ino))
  i = 1;
 else i = 0;
 unlink (filename);
 unlink (other);
 return (i);
}
```

Your existing program may not work when the system is converted to long file names. Software that opens directories, reads the directory entries directly, and expects the size of directory entries to be a constant. This software should be changed to use directory library routines or use `getdirentries` system calls. You should include `ndir.h` instead of `sys/dir.h` when you use the directory library.

This example works only with short file name systems.

```
/* This routine accept 2 arguments. The first is a directory path and
   the second is a file name which is to be searched in the directory.
   This routine will work in a short filename file system but NOT in a
   NFS environment nor in a long filename file system. */
#include <sys/dir.h>
find_name(dnamep, fnampe)
register char *fnamep, *dnamep;
{
   register int i, len;
   register int fd;
   struct direct ds;

   fd = open(dnamep, O_RDONLY));
   len = strlen(fnamep);

   while (read(fd, &ds, sizeof(struct direct)) == sizeof(struct direct))
      if (ds.d_namlen == len && !strcmp(ds.d_name, fnamep)) {
    printf("%s is found\n", fnamep);
    close(fd);
    return(0);
      }

   printf("%s is not found\n", fnamep);
   close(fd);
   return(1);
}
```

The examples continue on the next page.

This example works with short and long file names.

```
/* This routine accept 2 arguments. The first is a directory path and
   the second is a file name which is to be search in the directory.
   This routine will work in a NFS environment, long filename and
   short filename file systems.*/
#include <ndir.h>
find_name(dnamep, fnamep)
register char *fnamep, *dnamep;
{
 register int i, len;
 DIR *dirp;
 struct direct *dp;

 dirp = opendir(dnamep));
 len = strlen(fnamep);

 while (dp=readdir(dirp)){
    if (dp->d_namlen == len && !strcmp(dp->d_name, fnamep)) {
  printf("%s is found\n", fnamep);
  closedir(dirp);
  return(0);
    }
 }
 printf("%s is not found\n", fnamep);
 closedir(dirp);
 return(1);
}
```

The examples continue on the next page.

Software that assumes a maximum length of 14-character file names. For example:

```
char filename[14]
```

If you need only a few buffers, use `MAXNAMLEN` for buffer size. Otherwise, you need to allocate memory dynamically. If you need to allocate memory dynamically, you can obtain the file name size from `d_namlen` in the `struct direct`. For example:

```
/* This routine stores all file names in a directory in memory.*/
#include <ndir.h>

store_filenames(dnamep)
char *dnamep;
{
 DIR *dirp;
 struct direct *dp;
 char *cp;

 dirp = opendir(dnamep);
 while (dp = readdir(dirp)) {
  cp = malloc(dp->d_namlen + 1);
  strcpy(cp, dp->d_namlen);
   .
   .
   .
 }
 closedir(dirp);
}
```

The next page has the final example in the series.

Software may use `dirsiz` (assume it is a constant of 14) and use the value to mean the maximum length of the file name. For example:

```
char filename[DIRSIZ];
```

Use `MAXNAMLEN` as the maximum length of file names. Use `DIRSIZ_CONSTANT` as the maximum file length on a short file name system. If the `DIRSIZ_MACRO` compilation flag is turned on, `DIRSIZ` is a macro instead of a constant of 14. It accepts an argument that is the pointer to a `struct direct`. It returns the actual size of the directory entry in a long file name system. `DIRSIZ_MACRO` is mainly for porting programs that were originally developed under BSD 4.2 UNIX ™.

Software that uses `MAXNAMLEN` and assumes that `MAXNAMLEN` equals 14. `MAXNAMLEN` is now 255. Most of the software need only be recompiled, although the memory allocation could be fairly large.

Software that includes `dir.h` and uses `struct direct`. The `struct direct` for short file names is a fixed-size structure while the `struct direct` is variable length for long file names. HP recommends that software should be changed to include `ndir.h` and to use directory library routines.

Software that assumes there is only one file system magic number. The magic number for long file names is different than that for short file names. Software should be changed to allow the new magic number.

In general, all programs and routines that include `sys/dir.h` or use `FS_MAGIC` need to be reviewed and possibly changed to work with long file names.

Removing a Product from HP-UX

Besides the core filesets that constitute HP-UX, you can update your system by adding the filesets that make up a product. For example, setting aside any required customization or configuration, updating HP-UX so it has the SoftPC product amounts to adding one or more filesets to the existing HP-UX system. In time, you may decide you do not need a product. When this happens, you can remove the product and not otherwise affect the operation of HP-UX.

Prerequisites and Conditions

- Identify the product you want to remove according to the filesets it contains. You can see the filesets in `/etc/filesets`. You can verify the filesets in a product by looking in the documentation for the product. If you use the manual procedure, note the filesets that constitute a product so you can specify them to the `/etc/sysrm` command.
- The manual procedure cannot remove a fileset on a remote system (for example, an NFS-mounted file system).

SAM Procedure for Removing a Product (Subsystem)

While they may not technically be products, SAM lets you remove certain subsystems from HP-UX.

```
System Administration Manager
     ↓
Kernal Configuration ->
     ↓
Add/Remove Subsystems ...
```

At this point, you can remove the following subsystems:

- `NS/9000`
- `LAN/9000`
- `NFS/9000`
- `CR-ROM/9000`
- `DSKLESS/9000`

Follow the directions for removing a subsystem. If you need help, call the Help Screen.

Manual Method for Removing a Product

1. Log in as the root user.
2. Change to the root directory if necessary (`cd /`).
3. Shutdown the system to be in the single-user run level (`shutdown`).
4. Execute the following command where *product* is the name of a fileset (or filesets) determined in the prerequisite step:

   `/etc/sysrm` *product*

   For example, the following three commands remove the SoftPC product from your system:

   ```
   /etc/sysrm SPC_XWIN
   /etc/sysrm SPC_TERM
   /etc/sysrm SPC_CORE
   ```

   Later, you could load the product back onto the system by using the update procedure described in the earlier chapter named "Updating HP-UX".
5. After you remove the desired products, return to the multi-user run level by executing:

   ```
   reboot
   ```

Adding or Removing an NFS File System

In the earlier chapter named "Constructing an HP-UX System", you saw information about adding various networking systems. Besides going to other documentation to add or remove an NFS file system, you can perform this task with SAM.

Prerequisites and Conditions

The networking must operate properly. If you have `nfs` in the kernel, you must set up an NFS client according to the networking documentation.

SAM Procedure for Adding an NFS File System

```
System Administration Manager
    ↓
File Systems ->
    ↓
Add a Remote (NFS) File System ...
```

At this point, you see the data-entry screen. Enter appropriate data in each field. If you need help with with a particular field, call the Help Screen for that field.

SAM Procedure for Removing an NFS File System

```
System Administration Manager
    ↓
File Systems ->
    ↓
Remove a Remote (NFS) File System ...
```

Follow the directions for removing an NFS file system. Call the Help Screen as required.

Manual Procedure for Removing an NFS File System

See the manual named *Using and Administering NFS Services* for your series.

7

Managing the LP Spooler

The Line Printer (LP) Spooler is a set of utilities and commands that accepts requests to print files. The spooler subsequently schedules and processes the printing of the files. An initial module defines terms related to using the LP Spooler. The following table shows how subsequent modules can help you.

Name of Module	What The Module Does	SAM?
Setting up the LP Spooler	Describes the initial setup and how to get the spooler going.	No
Adding a Printer to the Spooler	Explains how to add a local and remote printer.	Yes
Checking LP Spooler Status	Explains commands used to monitor status.	Yes
Controlling Spooler Operation	Explains the enable/disable, remove, cancel, shut/start operations.	Limited Yes
Controlling Printer Priorities	Explains commands and procedures for controlling priorities for scheduling requests.	Yes
Displaying LP Spooler Performance Analysis	Shows data you can obtain that relates to performance.	No

LP Spooler Terminology

class
: A group of printers. When users make print requests, they can specify a particular class of printers as the destination. The first available printer in the class processes each request.

device
: The device file that designates a printer as an LP spooler device.

destination
: The name of an individual printer or the name of a class that contains related printers

fence
: The minimum priority of a spooled print job. The default value is zero.

interface
: The script or program that processes the LP requests for the LP spooler.

model
: Models are shell scripts that execute commands allowing the LP spooler's print job scheduler `lpsched` to communicate with the printers. The `/usr/spool/lp/model` directory contains the scripts.

request
: The ID numbers used to identify print requests. Each time a user exectutes `lp`, the LP spooler creates an ID number for the print request.

scheduler
: A process that routes requests from a print queue to the printing destination.

`/etc/mklp`
: A script you edit to add a printer to the spooler. It also contains lists of supported printers and their print options.

Setting Up the LP Spooler

The LP spooler comes with the HP-UX system, but you must set up the spooler before a printer can operate in a multiuser environment.

Prerequisites and Conditions

- If you installed the 7.0 release of HP-UX, the files for setting up the LP spooler reside in `/etc`. If you updated to the 7.0 release, `/etc/newconfig` contains files for the printer that were modified. Examine these files before you copy them to `/etc` to avoid inadvertently modifying your printer's behavior.
- Each installed printer must have a raw device file. If you use SAM to set up a printer, SAM makes the printer device files.
- There is no SAM procedure for setting up the LP spooler.

Manual Procedure for Setting UP the LP spooler

1. Log in as `root`.
2. Edit `/etc/passwd` so it has the following entry:

   ```
   lp::9:2::/usr/spool/lp:/bin/sh
   ```

   The above entry grants ownership of the LP spooler to the user named `lp`.

   If you want password protection for the user named `lp`, edit `/etc/passwd` so the example given above contains an asterisk (*) in the password field:

   ```
   lp:*:9:2::/usr/spool/lp:bin/sh
   ```

   The *passwd*(1) entry in the *HP-UX Reference* explains why you might want to do this.
3. Edit `/etc/group` so it has the following entry:

   ```
   bin::2:root,bin,lp
   ```

   The above entry provides group ownership of the LP spooler to the user named `bin`.

4. Edit the `/etc/rc` file so it provides for local and possibly remote printers.

   a. For local spooling, make sure `/etc/rc` has these entries so the scheduler starts up each time you boot HP-UX:

   ```
   # Start lp printer scheduler
        /usr/lib/lpshut>/dev/null 2>&1
        if [ -s /usr/spool/lp/pstatus ]
        then
            rm -f /usr/spool/lp/SCHEDLOCK
            /usr/lib/lpsched
            echo line printer spooler started
        fi
   ```

   These entries start the LP scheduler for local printers each time you start up the system.

   b. For remote spooling, make sure `/etc/rc` has the following entries so the `rlpdaemon` starts up when you boot HP-UX:

   ```
   /usr/lib/rlpdaemon -l
   ```

   To specify the use of *inetd*(1M):

   ```
   /usr/lib/rlpdaemon -i -l
   ```

5. To make sure the spooler stops when you shut down the system, edit `/etc/shutdown` so it has this entry:

   ```
   /usr/lib/lpshut
   ```

Adding a Printer to the LP Spooler

Setting up the LP spooler provides a vehicle for printing. But you still need to add printers to the LP spooler.

Prerequisites and Conditions

- To add a printer, use SAM, work manually, or use the `mklp` script.
- Each printer needs a character device file. For example:

```
crw-------  4 lp bin 21 0x070100 Aug 17 12:48 lp
```

- The following entries show how to make raw device files for local and remote printers at major number 21 on the HP-IB interface at select code 7 and bus address 1. The documentation for installing peripherals has additional information.

```
mknod /dev/lp 4 21 0x070100
mknod /dev/rlp 4 21 0x070100
```

- SAM makes a device file for you when you add a printer.

SAM Procedure for Adding a Printer

System Administration Manager
↓
Peripheral Devices ->
↓
Printers and Plotters ->
↓
Add a Local Printer ...
↓

At this point, you provide information for the following items:

- `Printer name`
- `Printer model/interface`
- `Printer device file name` This item asks you for the device interface, select code, and bus address or channel number.
- `Printer priority`
- `Make this the system default printer?`
- `Printer class`
- `Printer connected to a terminal?`

To add a remote printer, select `Add a Remote Printer ...` and respond to the following items:

- `Printer name`
- `Remote system name`
- `Remote printer name`
- `Remote cancel model`
- `Remote status model`
- `Make this the system default printer?`
- `Printer class`
- `Restrict cancel>`
- `Remote printer on a BSD system?`

Manual Procedure for Adding a Printer

1. Log in as superuser (`root`) and shut down the LP scheduler:

   ```
   /usr/lib/lpshut
   ```

2. Examine the scripts in `/usr/spool/lp/model`. Select the model script you want. Make necessary customizations. A listing might look like this:

   ```
   -rw-r--r--  1  lp    bin    4677  Aug 24 18:47  hp2564a
   -rw-r--r--  1  lp    bin    3208  Aug 21 13:22  hp2932a
   -rw-r--r--  1  lp    bin    5037  Aug 27 09:02  hp2276a
   -rw-rw-rw-  1  root  other  3044  Sep 21 08:44  printmod
   ```

   In the listing, some permissions, ownerships, and groups are incorrect. Ensure the model script (or scripts) you select have these attributes:

   a. A permission mode of 644 (-rw-r—r—). If necessary, execute:

      ```
      chmod 644 printmod
      ```

   b. Owned by `lp`. If necessary, execute:

      ```
      chown bin printmod
      ```

   c. In the group `bin`. If necessary, execute:

      ```
      chgrp lp printmod
      ```

3. Use the `lpadmin` to add the printer to the LP spooler. Repeat the command for each printer you add to the system.

   The procedure continues on the next page.

For example, if you have an HP 2934A printer that is accessed through the device file `/dev/lp`, you can use the following command line:

```
/usr/lib/lpadmin -plp -v/dev/lp -mhp2934a -cclass1 -g4
```

The items in the examples work as follows:

`-plp`	Names the printer `lp`.
`-v/dev/lp`	Specifies the full path name of the printer's special file, the physical destination.
`-mhp2934a`	Specifies the HP2934A printer from the `/usr/spool/lp/model` directory.
`-cclass1`	Specifies the class (`class1`) that the printer `lp` belongs to. Class is optional; printers do not have to belong to a class.
`-g4`	Sets the default priority for the printer. This parameter is optional; the default is zero.

4. If you want to define the printer fence that is the minimum priority for a spool file to be printed, use `lpfence`. For example:

```
/usr/lib/lpfence  lp     2
```

The items in the examples work as follows:

`/usr/lib/lpfence`	The command.
`lp`	The printer name.
`2`	The fence.

The procedure continues on the next page.

5. Execute `accept` to allow the printer to accept print requests:

   ```
   accept lp3

   destination "lp3" now accepting requests
   ```

6. Execute `enable` to enable the printer to print the requests:

   ```
   enable lp3

   printer "lp3" now enabled
   ```

7. Restart the LP scheduler:

   ```
   /usr/lib/lpsched
   ```

8. Verify that the scheduler is functioning:

   ```
   lpstat -t
   ```

   The display could show information like this:

   ```
   scheduler is running
   system default destination: lp
   device for lp: /dev/lp3
   device for lp3: /dev/ttyd2p0
   lp accepting requests since Aug 17 10:29
   lp3 accepting requests since Aug 18 14:47
   printer lp  is idle. enabled since Aug 17 10:45
           fence priority: 0
   printer lp3 is idle. enabled since Aug 18 15:11
           fence priority: 0
   ```

9. If the scheduler is not running, remove the file `SCHEDLOCK`:

   ```
   rm -f /usr/spool/lp/SCHEDLOCK
   ```

   Try again to specify the default printer. The `SCHEDLOCK` prevents the execution of more than one scheduler. The `lpshut` command automatically removes the SCHEDLOCK file when it terminates the LP scheduler.

   You may also need to remove the file named `FIFO` before the scheduler works properly. A FIFO is a named pipe created by `lpsched` for LP scheduler communications. Remove this file by executing:

   ```
   rm -f /usr/spool/lp/FIFO
   ```

The mklp Script Procedure for Adding a Printer

You can edit and execute the `/etc/mklp` script to add a printer to the LP spooler.

The `/etc/mklp` script contains the following sections:

1. Section 1 sets the ownership, group and access mode for each file used by the LP spooler.
2. Section 2 lists each supported printer and lists the commands required to configure it.
3. Section 3 explains how to set up a default printer.
4. Section 4 describes how to configure a system for remote spooling.

Using the following procedure, edit `/etc/mklp` so it fits your system needs. Then, execute the script.

1. Create a device file. If necessary, see the chapter named "Managing Devices".
2. Connect the printer. See the documentation that came with your printer.
3. These lines appear in the script's introduction. Comment them out:

   ```
   echo "mklp:  template version -- customize script before using it"
   exit 1
   ```

4. In Section 2, uncomment the lines:

   ```
   cd /dev
   name= ?? (name of the device)
   dev= ?? (the device file)
   ```

The procedure continues on the next page.

5. Add the device name and the device file to the lines you just uncommented. For example:

```
cd /dev
name=lp (now contains lp, the name of the device)
dev= /dev/lp3 (now contains lp3, the name of
the device file)
```

6. Also in Section 2, select the device model by uncommenting a line that contains the model name. For example:

```
# model=hp2225a  # for HP 2225
  model=hp2563a  # for HP 2563
# model=hp2686a  # for HP 2686 (serial) laser jet
# model=hp2932a  # for HP 2932A
```

   In this example the line containing the HP 2563 printer is uncommented, so that printer model is selected.

7. Finally, in Section 2, uncomment the lines:

```
$lpshut
$lpadmin -p$name -v$dev -m$model -h
$accept $name
$enable $name
$lpsched
```

8. Uncomment the last line of the `mklp` script:

```
exit 0
```

Manual Procedure for Adding a Remote Printer to the LP Spooler

The LP spooler can schedule print jobs on printers that exist on remote systems.

Use the following syntax:

`/usr/lib/lpadmin -p`*name* `-v/dev/`*device* `-m`*model* `#` *remote_sp_options*

where the parameters work as follows:

name	The printer name that users specify when they execute `lp`.
device	Specifies the full pathname of the special (device) file of the printer.
model#	The printer model (must be listed in the `/usr/spool/lp/model` directory).
remote_sp_options	Any of the options shown in the table on the following page.

Option	Description
-ob3	Uses 3-digit print request numbers associated with the printer directory for compatibility with BSD systems. HP-UX uses 4-digit print request numbers.
-oci*rcancel*	Causes *cancel* to use *rcancel* to cancel requests to remote printers. Specify the full pathname to ensure that the correct command is used.
-ocm*rcancel*	The model *rcancel* cancels requests to remote printers.
-orm*system*	The name of the remote system is *system*.
-orp*printer*	The name of the printer to use on the remote machine is *printer*.
-orc	Restricts users to canceling only their own requests. The default is to not restrict the *cancel* command.
-osi*rlpstat*	Causes `lpstat` to use `rlpstat` to obtain the status of requests on remote printers. Specify the full pathname to ensure that the correct command is used.
-osm*rlpstat*	The `rlpstat` model obtains the status of requests to remote printers.

1. Here is an example of a command that enables remote spooling (with various options). Enter the command with no (Return) or use a backslash before any returns (as shown):

```
/usr/lib/lpadmin -plp3 -mrmodel -v/dev/null -ocm rcmodel\
   -osm rsmodel -ormsystem2 -orplp3 -ob3
```

In this example, the local printer is `lp3` and its model is called `rmodel`. The device destination is specified as `/dev/null` because the spooler requires a device name. In this case, networking software takes care of which device to send a request to.

2. On the remote system, add this line to the `/etc/rc` file:

```
/usr/lib/rlpdaemon -l
```

Or, in the `/etc/inetd.conf` file, add this line:

```
printer stream tcp nowait root /usr/lib/rlpdaemon rlpdaemon -i -l
```

Then execute the command:

```
inetd -c
```

3. Change the ownership of files to `root` as follows:

```
chown root /usr/lib/rcancel
chown root /usr/lib/rlp
chown root /usr/lib/rlpdaemon
chown root /usr/lib/rlpstat
```

Checking LP Spooler Status

You typically check the status of the LP spooler before you do anything to it. To do this, you can use SAM or execute certain commands.

SAM Procedure for Checking the LP Spooler

To view status information:

`System Administration Manager`

↓

`Peripheral Devices ->`

↓

`Printers and Plotters ->`

↓

`View Printer Status Information`

To view print requests, proceed as above, except select `View Print Requests`

Manual Procedure for Checking the LP Spooler With lpstat

Use the `lpstat` command to check the status of printers, print jobs or the scheduler. The `-t` option, provides complete LP spooler status information:

```
lpstat -t

scheduler is running
system default destination: lp
device for lp: /dev/lp
lp accepting requests since Jun 14, 15:37
printer lp now printing lp-58.  enabled since Jun 23 13:31
        fence priority : 0
lp-58           williams      priority 3 Jul  9  12:53 on lp
        services                              751 bytes
lp-59           jones         priority 1 Jul  9  13:39
        (standard input)                     3264 bytes
```

See the *lpstat*(1) entry in the *HP-UX Reference* manual to get information.

Controlling Spooler Operation

The scheduler routes `lp` requests to an interface program that controls printing on a printer or class of printers. If you stop the scheduler, it can no longer route print requests and the LP spooler stops working. Besides starting or stopping the lp scheduler, you can enable or disable printers and cancel print requests. This module explains how to perform these tasks.

Prerequisites and Conditions

- You must stop the LP scheduler when adding a new printer on the system or when moving requests from one printer to another.
- Once you have completed these tasks, start the LP spooler again.
- Depending on what you want to accomplish on a global basis, you coordinate starting, stopping, enabling, disabling, and canceling operations. While this manual describes tasks in separate sections, you typically perform a series of tasks related to solving a global problem such as getting local and remote printers to function as desired.
- SAM performs most operations related to controlling spooler operation. It does not let you cancel a print request.

SAM Procedures for Controlling Spooler Operation

`System Administration Manager`

↓

`Peripheral Devices ->`

↓

`Printers and Plotters ->`

↓

At this point, you can select the following options:

- `Enable a Printer ...` Asks for the printer name.
- `Disable a Printer ...` Asks for the printer name.
- `Remove a Printer ...` Asks for the printer name.
- `Shut down/Start up the Spool System`
- `Set the System Default Printer` Asks for the printer name.

Manual Procedures for Controlling Spooler Operation

The following items explain the operations. Find the one you need and perform the task.

Starting the LP Scheduler

```
/usr/lib/lpsched
```

Stopping the Spooler

```
/usr/lib/lpshut
```

Moving Print Requests

HP does not recommend using `lpmove` to move requests associated with remote printers. The remote IDs may conflict and destroy a request.

1. Verify that a device is accepting print requests before moving the print requests to that device:

   ```
   lpstat -alp2

   lp2 accepting requests since Fri Jun 10 21:42:35
   ```

2. Stop the LP spooler scheduler:

   ```
   /usr/lib/lpshut
   ```

3. You can use `lpmove` in one of the following ways:

   a. Move all requests for printer `lp1` to printer `lp2`:

      ```
      /usr/lib/lpmove lp1 lp2
      ```

   b. Move an individual request using its ID number `lp1-103` to printer `lp2`:

      ```
      /usr/lib/lpmove lp1-103 lp2
      ```

Canceling Local Print Requests

The procedure assumes a user named **ged**.

1. Use the `lpstat` command to get the request ID number:

```
lpstat -t

scheduler is running
system default destination: lp1
device for lp1: /dev/lp
lp1 accepting requests since Jun 14, 15:37
printer lp1 now printing lp-58.  enabled since Jun 23 13:31
        fence priority : 0
printer lp2 now printing lp-58.  enabled since Jun 23 13:31
        fence priority : 0
lp1-1207         bill        priority 17 Aug  10:15 on lp1
        applications                       7084 bytes
lp2-1212         ged         priority 17 Aug  10:24 on lp2
        (standard input)                  41211 bytes
```

2. Execute `cancel` with the request ID of the request you want to cancel:

```
cancel lp2-1212
```

   The system displays:

```
request "lp2-1212" canceled
```

Canceling a Remote Print Request

1. Execute the `lpstat` command with the `-d` (device) and `-u` (user) options to get the request ID number:

```
lpstat -dlj -ucath

scheduler is running
system default destination: lj
device for lj: /dev/lp
lj accepting requests since Jun 14, 15:37
printer lj now printing lp-58.  enabled since Jun 23 13:31
        fence priority : 0
lj-902        cath        priority 17 Aug  10:21 on lj
        applications                       9011 bytes
```

2. Execute `rcancel` with the request ID of the request you want to cancel:

```
rcancel lj-902
```

The system displays:

```
request "lj-902" canceled
```

Removing a Printer

Before you remove a printer from the LP spooler, you must either move or cancel that printer's print requests. These operations were shown earlier.

- To remove a printer from a specific class:

  ```
  /usr/lib/lpadmin -plp -rclass1
  ```

  When you remove all printers from a class, the class no longer exists.

- To remove an entire class of printers:

  ```
  /usr/lib/lpadmin -xclass1
  ```

- To remove a printer that is not a member of a class:

  ```
  /usr/lib/lpadmin -xlp
  ```

Preventing Acceptance of Requests

Use the `accept` and `reject` commands.

- To disable a printer:

  ```
  reject lp3
  ```

- To enable a printer:

  ```
  accept lp3
  ```

- To disable a class of printers:

  ```
  /usr/lib/reject class1
  ```

- To enable a class of printers:

  ```
  /usr/lib/accept class1
  ```

Setting the Default Printer

The default printer receives print requests unless you specify another printer. You can specify a system-wide default printer, or specify an individual default printer for each user.

You can set (or change) the default printing destination for all users using the `lpadmin` command. For example, to set `lp1` as the default printer for all users on a system:

```
lpadmin -dlp1
```

You can modify the message of the day file (`/etc/motd`) to contain the name and location of the default printer. When users log into the system, this information is part of the login message.

Using SAM to set the system default printer is part of the procedure for adding a printer. Also, you can use the `Set the System Default Printer` option.

You can specify default printing for each user (that is, a default destination) using an environment variable in the user's `.profile`, `.login` or `.cshrc` file.

If the user runs the Bourne shell (`sh`) or the Korn shell (`ksh`), add the following line to `.profile` in the user's home directory:

```
LPDEST=destination;export LPDEST
```

If the user runs the C shell (`csh`), add the following line to `.login` or `.cshrc` in the user's home directory:

```
setenv LPDEST destination
```

Controlling Printer Priority

The printing fence is the minimum priority for a job. Print jobs with a priority greater than the printing fence will print before those with a lesser priority. For example, if the printing fence is 2, print jobs with a priority of 3 or 4 print before those with a priority less than 2.

In SAM, you can set the printer priority in the screen for `Add a Local Printer`.

If you work manually and use the `lpfence` command to change the printing priority, proceed as follows:

1. Stop the spooler:

   ```
   lpshut
   ```

2. Use `lpfence` to change the printing priority:

   ```
   lpfence lp3 2
   ```

3. Restart the spooler:

   ```
   lpsched
   ```

In the example, the printing priority of `lp3` has been changed to 2.

Changing the Default Priority

Use the `lpadmin` command to change the default priority for a printer:

```
lpshut
lpadmin -plp3 -g4
lpsched
```

In this example, the default priority of `lp3` has been changed to 4.

Displaying LP Spooler Performance Analysis

The file `/usr/spool/lp/lpana.log` accumulates LP spooler performance information as long as the LP spooler scheduler is running. Use the `lpana` command to display an analysis of the LP spooler's performance. Enter `lpana` with no arguments to see a performance analysis of all printers and classes:

```
lpana
```

You see information like this:

```
performance analysis is done from Apr. 22 '89 through Jul. 26 '89
---printers ----wait----   ---print---   ---bytes--- -sum-   num_of
/classes      AV      SD      AV      SD     AV      SD     KB   requests
asp          5'00      1   30'00       9   3940    1213   2017        512
lp1         24'00     10    1'14      24   9527    5231    219         23
lp2          1'37     20    5'01      12    825     917   3495       4236
/one        23'45      5   58'12      33   8873    4798   5847        659
/two         7'12      2   19'01       4    630     212   2591       4112
```

Classes are designated by a slash (/) before the class name.

Wait AV	Average time a print request waits in the queue before printing.
Wait SD	Standard deviation for waiting time.
Print AV	Average total printing time for a print request.
Print SD	Standard deviation for printing time.
Bytes Av	Average number of bytes per print request.
Bytes SD	Standard deviation of the number of bytes.
Sum KB	Sum of bytes for all print requests (in Kbytes).
Num of Requests	Total number of print requests logged since LP spooler was last started.

To display performance analysis information on a particular printer or class of printers, use the `lpana` command with the `-d` option and the printer or class name. For example, to display information on the printer `lp1`:

```
lpana -dlp1

performance analysis is done from April 22 through July 26
printers  Wait        Print      Bytes          Sum    Num_of
/classes  AV     SD   AV    SD   AV     SD      KB     Requests
lp1       24     10   1'14  24   9527   5231    219    23
```

LP Spooler Directories

`/usr/spool/lp`	LP spooler system parent directory. All information about the setup and printing queues is located here.
`/usr/spool/lp/class`	Printer classes directory. This contains the files that define how printers are grouped.
`/usr/spool/lp/model`	System-supplied interface programs. This directory contains the model shell scripts designed for particular printer models.
`/usr/spool/lp/interface`	Interface programs in use on your system. This has shell scripts from `/usr/spool/lp/model` that may be modified for particular printers. If interfacing a printer for which there is no model file, you may need to create an interface program for it.
`/usr/spool/lp/request`	Destination queues. This is where all `lp` requests are queued. It usually contains a subdirectory for each printer configured on the system.
`/usr/bin`	Contains user-executable commands, such as the LP spooler commands that general users can execute.

`/usr/lib`	Contains administrator-executable commands, such as the LP spooler commands that only `root` or `lp` can execute.
`/usr/spool/lp/cmodel`	Contains system-supplied interface programs, in the form of model shell scripts for processing remote cancel requests.
`/usr/spool/lp/smodel`	Contains system-supplied interface programs, in the form of model shell scripts for processing remote status requests.
`/usr/spool/lp/member`	Lists all printers, one file per printer.
`/usr/spool/lp/cinterface`	Contains the shell scripts from `/usr/spool/lp/cmodel`.
`/usr/spool/lp/sinterface`	Contains the shell scripts from `/usr/spool/lp/smodel` for the installed printers.
`/usr/spool/lp/fonts`	Contains fonts for LaserJet printers.

8

Managing Devices

In this manual, a device refers to any functional hardware product you connect to the system (mouse, keyboard, printer, monitor, plotter, modem, bar-code reader, and so on). This definition anticipates that you will determine the required interface cards, cables, connectors, device protocols, and standards as a part of managing the devices.

This chapter has the following purposes:

- To provide an overall view of managing devices. This view gives attention to configuration, installation, kernel drivers, and customization.
- To indicate how to make device files for commonly used devices.
- To show how SAM can manage devices.
- To point you to documentation related to devices.

The chapter does not provide a comprehensive picture of managing devices. To some extent, you must coordinate several sources of information. Besides getting appropriate information, you typically perform tasks such as installing the device, reconfiguring the kernel, adding software to your system, customizing your system, and making device files.

You often need to solve problems related to fitting a device into your system so it does not conflict with existing devices. At the least, you need to coordinate the major numbers, select codes, unit numbers, bus addresses, and types of your devices. Within your total system, each device must have its place and not conflict with another device.

An Overview of Managing Devices

The following sections explain the tasks you perform.

Configure the Device

Make sure the device will function with your particular system. Make sure HP supports the device, or know how you can accommodate the device. Make sure you have the necessary connectors, cables, interface cards, adapters, and accessories. This manual does not provide much information about this. To get information about configuring a device, see your HP Sales Representative. Do not begin installation until you have accounted for every component part of a device.

Install and Test the Device

Shut down HP-UX and turn off existing devices. Physically connect the device to the HP-UX system. Set any switches so the device will not conflict with existing devices. Test the device. Do not start up HP-UX until you know the device functions properly. Use the documentation that came with the device. For many devices, you can get information about installation in the document for installing peripherals for your series.

The Kernel Needs a Device Driver

Each class of devices on the system must have a device driver. For example, if you add an HP 650A optical disk drive to your system (and any other device that uses a SCSI adapter), your kernel must have the `scsi` device driver. To check this, examine your configuration description file (`/etc/conf/dfile`). For example, try the following command:

```
grep sc /etc/conf/dfile
```

If you get `scsi`, you have the driver. If you get `*scsi`, the driver is commented out, and you will need to add it. If you get nothing, you need to add it. The chapter named "Reconfiguring the Kernel" explains this.

Each Device Needs a Device File

A device can have a block or character (raw) device file. It may need both types. The `/dev` directory (or one of its subdirectories) contains the block and character (raw) device files. In general, use a block device file for disks and tapes; use a character device file for other devices. You may also want character device files for disks or tapes other than the disk that contains the root file system. For example, if you add an HP 650A optical disk as a raw device with major number 47 at select code 17 and SCSI address 7, a listing of `/dev` should have a line like this:

```
crw-rw-rw- 1 root  other  47 0x110007 Aug 1 21:15 /dev/rscsi
```

If `/dev` does not have the device file, you can make it by executing:

```
mknod /dev/rscsi c 47 0x110007
```

In this command line: **/dev/rscsi** is the device file name, **c** tells `mknod` to create a character device file instead of a block device file, **47** is the major number for SCSI devices, **0x11** represents hexadecimal 17 for the select code, **00** is a bus address of 0, and **07** is the unit number (which is the SCSI adapter card address of 7 in this case). You get values such as select code and major number by examining:

- The document that came with your device.
- The document for installing peripherals for your computer series.
- The *mknod*(1m) and *mkdev*(1m) entries in the *HP-UX Reference*.
- The `/etc/mkdev` script.

You May Need to Update or Customize HP-UX

In some cases, you may need to add software to your HP-UX system to make a device function properly. In most cases, the software and documentation are included with the device. The chapter named "Updating HP-UX" in this manual explains the update process. Finally, you may need to customize certain files or alter the behavior of a script or utility (for example, a backup script or the line printer scheduler).

The Device Management Process

The previous section provided an overall view. This section describes the tasks you perform to make a device function with your system.

Prerequisites and Conditions

- Not setting up a device according to its documentation can be hazardous, expensive, and time consuming.
- When you install an entire system, plan the installation of devices so they work appropriately (for example, plan for switch settings, select codes, bus addresses, ports, protocols, major and minor numbers, and so on).
- When you add a device to an existing system, find an "opening" for the new device. For example, you may have devices assigned to bus addresses 0, 1, and 2 on an HP-IB port. This means a new device could use addresses from 3 to 7.
- Coordinate the devices, device drivers, and device files (both block and character(raw) files). The file named `/etc/conf/dfile` has the device drivers and the directory named `/dev` (or one of its subdirectories) has the device files.

SAM Procedure for Managing Devices

`System Administration Manager`

↓

`Peripheral Devices ->`

↓

The options continue on the next page.

In this menu, you have the following non-printer options:

- `Add a Terminal or Modem ...`
- `Add a Hard Disk Drive ...`
- `Remove a Hard Disk Drive ...`
- `Change a Hard Disk Drive Address ...`
- `View Disk Space Information`

Still in the same menu, you have the following sets of options for `Printers and Plotters ->`:

- For printers and plotters, you have these options:
  - `Add a Local Printer ...`
  - `Add a Remote Printer ...`
  - `Enable a Printer ...`
  - `Disable a Printer ...`
  - `Remove a Printer ...`
- For LP spooler administration, you have these options:
  - `Set the System Default Printer ..`
  - `Shut Down/Start Up the Spool System`
  - `View Printer Status Information`
  - `View Print Requests`

A General Manual Procedure for Managing Devices

Look for device drivers in `/etc/conf/dfile`. If you must add a driver, do so after you finish the tasks in the following list and before you return to the normal run-level. The chapter named "Reconfiguring the Kernel" explains how to add a driver.

1. If you attempt to install an unsupported device, see your HP representative to determine the implications.
2. Unpack the device. Read directions, notes, warnings, and cautions. Check the components against the invoice.
3. Read the documentation for the device. The documentation for the device may explain the installation. Note any special information(model number, serial number, switch settings, address, select code, type file, port, protocol) so you can use it in the next step.
4. Look up your device by model in the document for installing peripherals for your system. This document often provides enough information to install the device. For some devices, you may need to use other documentation.
5. Get your HP-UX system into the correct run-level according to the documentation for the device. This varies from multi-user state, single-user state, to completely shut down. If you install a card, shut down HP-UX and turn the computer (and expander) OFF.
6. Install the device according to its documentation, the information found in the document for installing peripherals for your series, and any other sources related to the device. For example, to install a SCSI device, you may need to examine several documents.
7. Test the device according to its documentation. This helps you avoid problems during start up and customization.
8. Start up HP-UX, if necessary, and make the device file for the device. The table on the following page shows how to make device files for selected devices.

To Make This Device File ...	Execute This Command ...
HP 7958A root drive at select code 14	`mknod /dev/dsk/0s0 b 0 0x0e0000`
Add a second similar disk used as a non-start up drive	`mknod /dev/rdsk/1s0 c 4 0x0e0100` `mknod /dev/dsk/1s0 b 0 0x0e0100`
HP 9144 cartridge tape drive at select code 7 and bus address 7	`mknod /dev/rct c 4 0x070700`
HP 7978A 9-track tape drive at 6250 bpi, Berkley auto-rewind, select code 14, and bus address 3 (make the `/dev/rmt` directory if necessary)	`mknod /dev/rmt/3s0 c 9 0x0e0382`
HP 2686A and HP 33440A Laserjet printer at select code 10 and port 2	`mknod /dev/lp2686 c 1 0x0a0204`
HP 650A optical disk (SCSI) at select code 17, bus address 0, and SCSI address 7 (the unit number)	`mknod /dev/rscsi c 47 0x110007`
Terminal on built-in interface at select code 9	`mknod /dev/tty02 c 1 0x090004`

Hayes Smartmodem 1200	`mknod /dev/tty02 c 1 0x090000`
	`mknod /dev/cua02 c 1 0x090001`
	`mknod /dev/cul02 c 1 0x090001`

HP 46084A ID module as the 2nd device in the HP-HIL interface	`mknod /dev/security c 24 0x000020`
HP 46087A digitizer as the 3rd device on the HP-HIL interface	`mknod /dev/digitizer c 24 0x000030`

9

Reconfiguring the Kernel

The installation process builds a kernel that has built in performance features. The kernel resides in the root file system as `/hp-ux`.

This chapter explains how to change the following things:

- Adding a device driver.
- Altering swap space.
- Changing the values of operating system parameters (kernel parameters).

You may need to change these things, for example, if your current need for swap space exceeds the amount you originally established.

In general, you seldom need to customize the kernel, especially if you install a system that meets your immediate and anticipated needs.

A newly created kernel replaces the existing kernel. The kernel is always built from scratch. You cannot add anything to an existing kernel.

To provide a convenient base for presenting information, the chapter discusses altering the kernel in terms of editing a configuration description file named `dfile` and subsequently using the `/etc/config` command to build the new kernel. Be aware, however, that SAM provides the capability to reconfigure the kernel. As before, it provides the preferred way to perform system administration tasks.

Planning to Reconfigure the Kernel

You can reconfigure the kernel over a range of complexity. For simple changes, you can use SAM or `/etc/config` to do the work. For complex changes, you might need to rewrite source code and add it to the kernel.

As you increase your understanding of HP-UX and your expertise in using the C programming language, you can work at more complex levels. You might eventually reach a level at which you alter the kernel in ways that Hewlett-Packard Company does not support (for example, developing a driver for a disk drive not supported in `/etc/disktab`). In these areas, you are on your own.

Identifying Things You Can or Should Reconfigure

Staying within supported boundaries, the following items describe things you can reconfigure:

- The device drivers that the kernel uses to work with peripheral devices, interface cards, and certain software applications. To do this, you typically:
  1. Make sure the files for the device drivers are installed on the system; and
  2. Edit the configuration description file, which is named `dfile`.
- The operating system uses swap space when it executes processes. The installation process creates swap space, but over time, you may need to add swap space to run the system. You can reconfigure this space, and chapter 6 describes some ways to do this. This chapter describes how swap space relates to the configuration description file.
- The parameters in your kernel have default values. These values work for most situations. You may need to change certain operating system parameters to make your software work.

Making Decisions about Reconfiguration

Know what you want to reconfigure (a device file, an application, the swap space, and so on). This need grows out of how people use the system.

Know your limits. Do not try to write a device driver and add it to the kernel if you have no expertise in writing drivers. You might get the system into a state in which you have no option but to reinstall it.

Match the use of commands to the task you want to perform (for example, decide to use `/etc/config` or `/usr/bin/sam`.

Consider the merit of running a default system that works adequately against altering your system until you get maximum performance. The system should work according to your needs. On the other hand, by creating a non-default system, you can introduce conditions that make the system unstable or inefficient.

Using Appropriate Commands to Reconfigure the Kernel

This chapter anticipates that you will use SAM to reconfigure the kernel. In this context, you may need little additional information because SAM provides online help as you move through a task. The manual method for reconfiguring the kernel uses the `/etc/config` command after you edit a configuration description file in which you can customize the following things:

- Add kernel device drivers.
- Alter swap space.
- Alter values of operating system parameters.

If you use the manual method, you may need to work through most of this chapter.

Be aware that with the 7.0 release, you cannot use the `/etc/reconfig` command. For administrators who have used pre-7.0 releases of HP-UX, the last module describes relationships between the now obsolete facility named `/etc/reconfig` and the newly implemented `/usr/bin/sam` (the System Administration Manager, or SAM).

Access to Source Code

Under certain agreements with Hewlett-Packard Company, you may obtain access to source code. In this capacity, you may be able to alter the kernel itself, provided you have the necessary expertise.

This manual does not discuss how to edit or develop source code that lets you reconfigure the kernel. You may find some useful information about this in the *HP-UX Driver Development Guide,* and perhaps in the *HP-UX SCSI Technical Reference Manual* or the *HP-UX Portability Guide.*

Files Related to Kernel Reconfiguration

/usr/bin/sam	A menu-oriented facility that leads you to data-entry screens, where you can perform tasks that reconfigure the kernel. SAM can also perform other system administration tasks.
/etc/config	A program used exclusively to reconfigure the kernel after you have prepared a configuration description file named `dfile`.

You also need the following files on the system:

/etc/master	A file that contains information used for generating a kernel.
/etc/conf/*.a	Libraries containing kernel code.
/etc/conf/h/*.h	Header files.
/etc/conf/machine/*.h	More header files.
/etc/conf/dfile.*	The sample configuration description files you can use as your `dfile` (for example, `dfile.min` or `dfile.full.lan`).
/etc/conf/dfile	The configuration description file for your kernel.

If the system does not have these files, load the KERN_BLD fileset from your original installation media. See "Updating HP-UX", if you need to add this fileset. Listing `/etc/filesets` should show KERN_BLD.

Using SAM to Reconfigure the Kernel

The following series of SAM menus shows the major ways SAM can reconfigure the kernel.

You can configure the device drivers.

`System Administration Manager`
↓
`Kernel Configuration`
↓
`Change I/O Configuration`
↓

The last menu above will shows several options for drivers (for example, disk, tape, printer). Select the one you need and enter the data.

You can modify operating system parameters.

`System Administration Manager`
↓
`Kernel Configuration`
↓
`Modify Operating System Parameters`
↓

The last menu above will show several options for modifying parameters (for example, Message, Semaphore, Accounting). Select the one you need and enter the data.

You can add to or delete from the swap space.

`System Administration Manager`
↓
`Kernel Configuration`
↓

At this point, you pick one of the following options and work through its data-entry screen:

- `Add to the Swap Configuration`
- `Delete from the Swap Configuration`

Also, you can add swap space if you use SAM to add a hard disk drive.

Other Kernel Reconfiguration Tasks Done by SAM

Besides the three tasks just shown, SAM lets you add or remove certain subsystems and generate a new kernel. To do this select `Kernel Configuration`, and then select among the following options:

- `Add/Remove Subsystems (NFS, LAN, NS, CD-ROM, DSKLESS)`
- `Generate a New Kernel (optionally reboot)`

Creating the Configuration Description File

If you use `config`, you first create a configuration description file named `/etc/conf/dfile` and then run `config`. If you use SAM, the help screens should provide the information you need. If they do not, read this chapter to get an overall picture.

Prerequisites and Conditions

- Name the configuration description file `dfile` unless you know why it must have another name.
- The dfile is a CDF in an HP-UX cluster; you *must* name it `dfile`.
- The upcoming procedures assume you copy a sample dfile and edit it.
- If you make major changes in the kernel, selecting a sample dfile can decrease the required editing. For minor changes, it is easier to edit your existing dfile.
- To create the new dfile, you have the following options:
  1. Edit your existing (`/etc/conf/dfile`).
  2. Copy a sample configuration description file (for example, `/etc/conf/dfile.full.lan`) to `/etc/conf/dfile` and edit it.

- In the dfile, you can select device drivers, specify alternate swap space, and/or assign alternate values to parameters. The following example shows parts of the file:

```
* HPUX_ID: @(#)dfile.full.lan 555.1   date
 ...
* ... You can ... using SAM or config(1m)
 ...
* DEVICE DRIVERS
cs80
 ...
* Networking drivers
lla
lan01
 ...
* SWAP CONFIGURATION
*
* CONFIGURABLE PARAMETERS
```

Manual Procedure for Creating the dfile

Creating a configuration description file occurs in four steps (described in subsequent sections). Work through only those steps in which you need to edit the file.

Step 1: Select a dfile.

Step 2: Add Device Drivers

Step 3: Specify the Swap Space

Step 4: Define Kernel Parameters.

Step 1: Select a dfile

1. You need a dfile that specifies what you want the kernel to provide. Your existing dfile may be adequate, except for minor editing. If it is, you need not copy a sample dfile. Move ahead to the next step according to your needs.

   If you plan to copy a sample dfile, the following table explains the files. Besides examining the table, you can see the sample files in `/etc/conf`. When you decide which sample file to use, *do NOT remove the files or alter them.* Instead copy your sample to your dfile by executing:

   `cp /etc/conf/dfile`*sample* `/etc/conf/dfile`

   where *sample* names the sample you want (for example, `.min`, `.full`, `.full.lan`).

2. When you have a dfile to work on, move ahead to the next step according to your needs.

Sample in /etc/conf	Description	Function
dfile.full	Configures HP-UX with all the supported device drivers in this release, minus LAN, NFS, cluster drivers, and pseudo drivers other than `pty`.	A fully loaded system less LAN
dfile.full.lan	`dfile.full` plus the LAN driver.	Default (and fully loaded) O.S. for a standalone system
dfile.min	Contains one of the minimal operating system configurations.	Minimal O.S. on a standalone system
dfile.maxservr	Contains the drivers in `dfile.full.lan`, plus the cluster drivers and operating system parameters required for a cluster server. The recommended dfile cluster server.	Default (and fully loaded) O.S. for cluster server
dfile.minservr	Contains all drivers in `dfile.min`, plus the cluster drivers and operating system parameters required for a cluster server.	Minimal O.S. for a cluster server
dfile.cnode	Contains a minimum kernel configuration for a client on an HP-UX cluster.	Default (and minimal) O.S. for client
dfile.cnodemax	Same as `dfile.maxservr` except replace cluster server-specific information with client-specific information.	Fully loaded O.S. for client
dfile.cnodemin	The minimum kernel configuration for a client (will not work on many clients; included only as a sample of a bare minimum kernel)	Illustration

Do not go past this section until you have selected your dfile.

Step 2: Adding Device Drivers

In Step 1, you created the configuration description file. Your kernel gets configured according to the device drivers specified in your dfile *each* time you create a new kernel.

The sample dfile you selected to be `dfile` is probably close to what you need, but it might not provide for all situations.

The kernel needs drivers for the following types of devices:

- Hardware such as disk drives, printers, plotters, scanners.
- Backplane cards such as HP-UX, VME, RS-232-C, SCSI.
- Software such as windows, LAN, clusters. Such drivers may be called pseudo drivers. They do not talk to hardware devices, but the kernel needs them.

Your system must have a correspondence among these items:

- The device drivers specified in `dfile`.
- The device files in `/dev`.
- The physical and logical devices attached to the system.

Your dfile must specify the types of device drivers your system will need after you configure the new kernel. Edit `dfile` (adding or removing device drivers) as required to specify the required drivers. Later, make sure `/dev` and its subdirectories have the related device files.

To provide the necessary device files, work through the following steps:

1. Study your system to see what devices it has according to the model numbers or names. Include peripheral devices, I/O cards, and software (LAN, X11 Windows, Windows/9000, RJE, SRM).

2. Examine `/etc/master` to see if HP supports the device. The following example shows part of the file. In your `dfile`, specify a device driver by using the device's handle.

```
name      handle       type    mask    block    char
----      ------       ----    ----    -----    -----
cs80      cs80          3      3FB      0        4
 ...
snalink   snalink       1      1C0     -1       36
 ...
dskless   dskless       18     100     -1       -1
 ...
console   cons          D      FD      -1        0
 ...
7958B     cs80
 ...
lp        printer
 ...
```

   The remaining parts of `/etc/master` have a similar appearance. If you cannot find a device in this file, HP-UX does not support the device. (This manual does not deal with unsupported devices).

   The steps continue on the following page.

3. Comment of uncomment the lines to add or delete device drivers. You need only one device driver per category of devices. For example, for five CS80 devices, you need only one `cs80` entry. The following table shows the major driver names (handles) and how the kernel uses them.
4. When you have added the device drivers, you have the following choices:
   a. If you need to alter the swap space, move ahead to "Step 3".
   b. If you do not need to alter the swap space, but do need to reset values of kernel parameters, move ahead to "Step 4".
   c. If you do not need to alter swap space or define parameters, move ahead to the module named "Using the config Command" and reconfigure the kernel.

Driver Name	Why the Kernel Uses It
cs80	Most mass storage devices (exists in all sample configuration files)
scsi	Small Computer System Interface direct access storage devices
amigo	Amigo mass storage devices
ciper	Ciper printers
printer	Non-ciper printers
hpib	Plotters, also include for Device I/O Library (DIL)
tape	9-Track magnetic tape drives
stape	9-Track streaming tape drives
ac	Optical autochanger
dos	HP 98686 DOS Coprocessor driver
vme	HP 98646 VME card
vme2	HP 98577A VME expander
nfs	Support for NFS networking
lla lan01	Support for NS-ARPA networking (formerly the ieee802 and ethernet drivers)
rfa	RFA server
dskless	Diskless code pseudo-driver
98624	Internal disk controller (always included by `config` because other drivers depend on it)
98625	Hi-speed HPIB disk controller
98628	RS-232 datacomm card
98642	RS-232 4-channel MUX card
98626	HP 98626 or HP 98644 RS-232 serial interface
gpio	GPIO card; also include for Device I/O Library (DIL)
srm	Shared Resource Management (SRM)
rje	Remote Job Execution (RJE)
ptymas ptyslv	Pseudo terminal drivers (required for HP Windows/9000, Xwindows, and other software). Included in all sample configuration files.

Step 3: Altering Swap Space

This section explains how to alter the swap space on your system in terms of reconfiguring the kernel. Before you begin the process, be aware that *altering the swap space on an existing system is NOT a trivial task.*

A system reserves a contiguous area (called swap space) on the secondary storage for use by the virtual memory management system. The virtual memory management system moves processes and shared memory objects to swap space when it has insufficient available memory. Your system created an amount of swap space during installation of HP-UX. The idea here is that you need to alter it.

Hewlett-Packard Company recommends using a hard disk drive for swap space even though, technically, you can swap to other media such as an optical disk.

One part of the task is editing your `dfile`. You are doing that part here. The steps you see later in this section mention other tasks. You must complete all the tasks related to adding swap space *BEFORE* you run a command that reconfigures the kernel. Reconfiguring the kernel may lead to rebooting the system. If the system tries to reboot and you did not provide for swap space, the reboot will fail.

Do not specify swap space in your `dfile` and forget the other tasks. Chapter 6 has information about adding swap space, installing another disk, initializing media, making a file system, mounting a file system, and using a second disk to hold part of the swap space. You must account for all of these things.

If you need more information about swap space before you decide to alter it, see the *HP-UX System Administration Concepts* manual.

You can have single or multiple swap disks. Use the following information to specify swap space on the swap disks:

1. In a default file system, one disk holds the file system *and* provides swap space. The kernel checks to see if the swap device has been changed; so to use non-default values for the swap devices, you must specify them in your configuration description file. The system uses this information to create the `conf.c` file. The process works as follows:

   a. If the swap device is specified in the configuration file, use the values from `conf.c`.

   b. If the swap device is not specified in the configuration file, HP-UX automatically assigns the root device as the swapping device.

2. Given the situation in Step 1, one way to increase the swap space is to add a second disk drive. The remaining procedure assumes both drives are CS80 devices and you use the default amounts of swap space. (Extrapolate from the examples for your case.) The *mkfs*(1M) entry in the *HP-UX Reference* manual and the `/etc/disktab` file contain information about this.

3. Determine the default amounts of swap space for your disk drive by examining information found in `/etc/disktab`. For example, `disktab` shows that an HP 7958 drive having 130.7 Mbytes defaults to 20 Mbytes of swap space as follows:

```
# HP7958 has 130.7 MBytes
#       256 Bytes/sector
#       63 sectors/track; 8 heads (small); 1013 cylinders;
hp7958:\
        :20 MBytes swap:ns#21:nt#6:nc#850:\
        :s0#107100:b0#8192:f)#1024:\
        :se#256:rm$3600:
```

The procedure continues on the next page.

4. Continuing the example, disk 1 is the root device and you enable it with a swap-auto entry. Disk 2 at Select Code 14 and Bus Address 1 is the second swap entry in `dfile`. You enter the following two lines after the commented line saying `* SWAP CONFIGURATION`:

```
swap auto   /* disk 1 */
swap cs80 E0100 -1   /* disk 2 */
```

   As another example, assume a root file system on an HP 7937 disk and you move `/users` onto a second, similar disk. You want both disks to provide default amounts of swap space. Both disks use Select Code 14 and have Bus Addresses of 0 and 1. They have device files named `/dev/dsk/0s0` and `/dev/dsk/1s0`. Add the following lines to `dfile`:

```
swap cs80 0e0000 -1    #rootdisk
swap cs80 0e0100  0    #swapdisk
```

5. Adding lines as shown helps your system provide the swap space. If you want the system to automatically access the new swap device during system startup, add an entry for the disk to `/etc/checklist` as shown in Chapter 6. Without having an entry in `checklist`, after you reboot the system, you must manually execute:

   `swapon` *2nd_disk_dev_file*

6. If you decide to add swap space this way, you must also install the second disk, initialize its media, make a file system on the disk, and provide for mounting the disk. See Chapter 6. If you add swap space this way, you should finish creating your `dfile`. But before you reconfigure the kernel, make a note to yourself to complete the tasks mentioned here if you have not previously done them.

7. After you edit your `dfile`, as indicated in step 6, proceed as follows:

   a. If you need to alter kernel parameters, move ahead to "Step 4: Altering Values of Kernel Parameters".

   b. If you do not alter parameters, move ahead to the module named "Using the config Command".

Step 4: Altering Values of Configurable Kernel Parameters

This is the last step in creating your `dfile`. Earlier, you may have added device drivers and/or added swap space. This final step lets you alter the values of kernel parameters.

Note The kernel HP-UX created during installation is designed to work well. In general, do not alter the values of kernel parameters unless you know why you want to alter them. Kernel drivers and operating system drivers mean the same thing in this chapter.

If you decide to alter some values, edit your `dfile`. If you need more information about kernel parameters, read the following material:

- The sections that appear after the following example.
- The appendix in this manual that deals with the kernel parameters.

The following example shows how you might alter kernel parameters on an HP-UX cluster server. You seldom need to alter parameters on a stand-alone workstation, and you may need to alter parameters on a multi-user system having multiple terminals.

```
* CONFIGURABLE PARAMETERS
maxuprc 64
msgtql 256
maxusers 16
num_cnodes 20
server_node 1
ncallout 984
parity_option 0
dskless_mbufs 31
num_lan_cards 4
netmemmax 500000
```

When you finish altering values of kernel parameters, move ahead to the module named "Using the config Command" and reconfigure your kernel.

The remainder of this step provides information if you need to use it.

Kernel Parameters You Can Alter in Step 4

The kernel parameters determine how the operating system manages memory, limits table sizes, and determines other operating system limits. Your HP-UX system has two types of parameters:

- Operating system parameters.
- System V IPC code capability parameters.

The following syntax shows the format of a parameter in your `dfile`:

`parameter` [*number* or *formula*]

where the parts work as follows:

number	A decimal number, or an octal number in C syntax.
formula	An arithmetic expression made up of numeric constants (in C syntax) and previously specified, tunable parameters. You cannot use names of parameters defined later in the dfile.

Using this syntax, the next several items indicate the types of kernel parameters you can redefine and show their names.

accounting code parameters	Used by system accounting. Parameters: `timeslice`, `acctsuspend`, `acctresume`.
time information	Used to determine the time from Greenwich Mean Time and differences due to daylight savings time. Parameters: `dst`, `timezone`.
parity errors	Action for a RAM parity error. Parameter: `parity_option`.
limiter for system resource allocation	Calculate values of global kernel parameters. Parameter: `maxusers`.
file system parameters	Number of open files, open inodes in the system, file system buffer headers, file locks, shared text descriptors (maximum), and file-size limit for processes. Parameters: `nfile`, `ninode`, `nbuf`, `nflocks`, `ntext`, `filesizelimit`.
process maximums	Maximum number of processes per user per system. Parameters: `maxuprc`, `nproc`.

maximum number of kernel timeouts	Maximum timeouts scheduled by the kernel at a time. Parameter: `ncallout`.
user process size limits	Maximum data, stack, and text size. See the later description "User Process Size Limits" below. Parameters: `maxdsiz`, `maxssiz`, `maxtsiz`.
memory parameters	Guarantee available memory for the virtual memory system and/or system overhead at any time and the area used by `argdev`. Parameters: `unlockable_mem`, `argdevnblk`, `dos_mem_byte`.
pseudo-teletypes	Determine maximum number of pseudo-teletypes. Parameter: `npty`.
number of DIL open device files	Parameter: `ndilbuffers`.
floating point accelerator capability	Parameter: `fpa`.
HP-UX cluster parameters	Configures arrays, swap and other HP-UX cluster structures. Parameters: `dskless_node`, `using_array_size`, `server_node`, `serving_array_size`, `maxswapchunks`, `minswapchunks`, `dskless_cbufs`, `dskless_mbufs`, `ngcsp`, `num_cnodes`, `dskless_fsbufs`,`selftest_period`.
total number of ITE text lines	Parameter: `scroll_lines`.
LAN parameters	Parameter: `num_lan_cards`, `netmeminit`, `netmemmax`, `netmemthresh`, `netisr_priority`.

The kernel parameters associated with process size interact with each other and need to be explained as a group.

The address space for a process consists of text space, data space, stack space, and possibly some shared memory segments. For example, `maxssiz` limits the stack size and will stop an infinite recursive program. The total process size is limited to 4 Gbytes, regardless of the size of these parameters. These parameters must work together, so be careful how you change them.

The HP-UX kernel has System V InterProcess Communication (IPC) code. It is not needed to run HP-UX, but your applications may use the System V IPC functions. Leave the System V IPC code in your kernel unless you are sure you don't need it.

The format for specifying System V IPC code in `config` is the same as for other kernel parameters.

System V IPC code has three parts: messages (mesg), semaphores (sema), and shared memory (shmem). Your kernel includes them unless you specifically exclude them with `mesg=0`, `sema=0`, and `shmem=0`.

For each of the System V IPC parts, there are several associated tunable parameters. If the part is set to 0, you cannot change any of its associated parameters. If the part is set to 1 (include in kernel), you may change any of its associated parameters.

mesg	Kernel code used for System V IPC messages. Parameters: `msgmap`, `msgmax`, `msgmnb`, `msgmni`, `msgseg`, `msgssz`, `msgtql`.
sema	Kernel code used for System V IPC semaphores. Parameters: `semaem`,`semmap`, `semmni`, `semmns`, `semmnu`, `semume`, `semvmx`.
shmem	Kernel code used for System V IPC shared memory. Parameters: `shmall`, `shmbrk`, `shmmax`, `shmmaxaddr`, `shmmin`, `shmmni`, `shmseg`.

Using the config Command

Once you edit your `dfile`, you can reconfigure the kernel.

Prerequisites and Conditions

- You worked through the process to set up `/etc/conf/dfile` so it has appropriate lines in it. If you edited the file extensively, you may have worked through many sections to get here, and you may have set up additional hardware. You need a configuration description file and appropriate hardware to use the `config` program.
- For an HP-UX cluster, you must reconfigure the kernel from the cluster server. You cannot use a cnode except to rebuild the kernel for that client.

Manual Procedure for Using config

1. You have some options concerning the kernel you intend to reconfigure:

   a. For a workstation or a multi-user system, log in as the root user. Tell users of terminals you will shut down the system. Then, shut down the system by executing:

   ```
   cd /
   shutdown 0
   ```

   b. For an HP-UX cluster, use the above procedure to reconfigure the kernel of the server after telling all clients you will shut down the system.

   c. For a client, log in as the superuser by executing the following command from the client whose kernel you plan to reconfigure:

   ```
   su root        Type the root password.
   password:
   ```

   (You could use `rlogin` from the cluster server to achieve the same result.) Do not shut down the system. Do not make a backup kernel. Move on to the step where you change to `/etc/conf`.

The procedure continues on the next page.

2. Except for a client in an HP-UX cluster, be in the root directory and back up the existing kernel by executing:

   ```
   cp /hp-ux /SYSBCKUP
   ```

   In an HP-UX cluster, `/hp-ux` is a CDF. Each client has a kernel under `/hp-ux+`. `/SYSBCKUP` is a file that must be usable by the cluster server. It is not a CDF. Therefore, do not perform this step if you are reconfiguring the kernel for a client.

3. Change to the directory having your configuration description file by executing:

   ```
   cd /etc/conf
   ```

   You can use a directory other than `/etc/conf`. But since all files required by `config` reside in `/etc/conf`, HP recommends using this directory. Do not execute `config` in the root directory or you will prematurely overwrite your kernel.

4. Execute `config` on your configuration description file.

   ```
   /etc/config dfile
   ```

   This creates `conf.c`, which is used to build your new kernel; and `config.mk`, which is a makefile.

   Use `ls` if you want to verify that the files were created.

5. Create the new hp-ux kernel in the current directory by executing:

   ```
   make -f config.mk
   ```

   During execution, `config.mk` lists portions of the kernel being built. For example, you might see:

   ```
   /etc/conf/libmin.a
      cs80.o
      mux.o
      muxs.o
      ⋮
   /etc/conf/libdevelop.a
      amigo.o
      ciper.o
      ⋮
   ```

6. Except for a client in an HP-UX cluster, move the new kernel to the root directory by executing:

   ```
   cp hp-ux /hp-ux
   ```

   This command also works if you are at a client and have made a kernel for that person's cnode. But if you are on the cluster server and just created a kernel for a client, you need the following syntax:

   `cp hp-ux /hp-ux+/`*client_name*

   where *client_name* might be `hpfcxyz`.

7. Reboot the system by executing:

   ```
   reboot
   ```

8. The system should start up and get into the multi-user state.

Possible Follow up Steps

1. If you are not running a server on an HP-UX cluster and if the newly configured kernel will not start up, use the boot ROM's attended mode to access the backup kernel saved during the process. To use the attended mode, press the space bar during bootup. This halts the automatic boot mechanism and allows you to choose the operating system to load. Select the option called `SYSBCKUP` rather than the new kernel `SYSHPUX`. Then, you can try work through the process again, attempting to determine the problem and solve it.
2. If you are in an HP-UX cluster at a client, do not boot using the backup kernel. Since the backup kernel applies to, the cluster server, you must instead create a new kernel from the cluster server, explicitly moving to the hidden directory in the CDF and working on the appropriate files.
3. If you added a second swap device, you must enable the disk before the system will use it. To enable the second device, execute:

   swapon *2nd_swap_dev_file*

   To enable the disk each time you start up the system, add an entry for the new swap device to `/etc/checklist`. This way, the `/etc/rc` script executes `swapon -a` and enable all devices specified in `/etc/checklist`.

10

Managing an HP-UX Cluster

This chapter explains how to set up and manage **HP-UX clusters.**

What is an HP-UX Cluster?

An HP-UX cluster is a network of HP 9000 computers, connected by LAN (Local Area Network) hardware and software, in which only one computer has file-system disk drives attached to it. This computer is known as the **cluster root server** (usually shortened to **cluster server** and occasionally to **server**).

The cluster server processes file system requests from the other machines in the cluster, which have no file-system disks and are called **cluster clients,** or occasionally **clients.**

The cluster server can also function as a workstation or multi-user system: it is not restricted to servicing file-system requests.

Each machine in a cluster (the cluster server or any client) can be referred to as a **cluster node,** sometimes shortened to **cnode.**

It is possible for cluster clients to have their own disk drives, but these can be used only for swap space, and even this is not the normal case: usually a cluster client will swap to the server's swap space.

The illustration on the next page shows how a cluster differs from a standalone system.

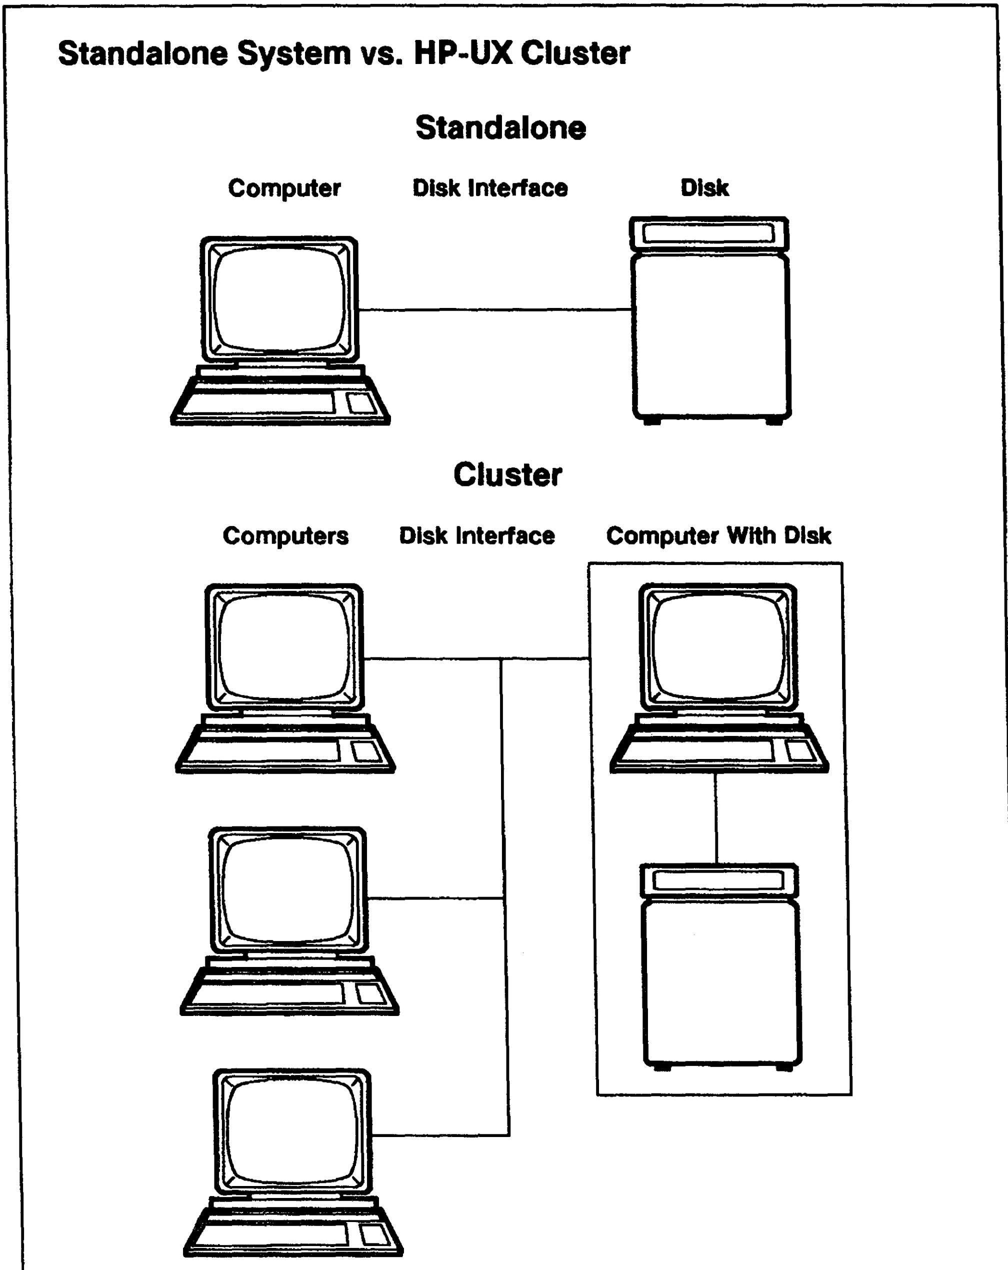

LG200151_001

What This Chapter Covers:

- Creating a cluster
  - Configuring the cluster server
  - Adding cluster clients
- Connecting your cluster to another network
- Removing cluster clients
- Administering an HP-UX cluster
  - Booting a cluster
  - Shutting down a cluster
  - Reconfiguring the kernel for a cluster node
  - Determining swap space
  - Adding a Local Swap Device
  - Installing and removing programs in a mixed cluster

Note If you are about to build a cluster for the first time, or have not managed a cluster before and now need to, you should read all of the Cluster Concepts chapter in *HP-UX System Administration Concepts* before you go any further.

What This Chapter Does Not Cover:

- Installing and configuring networking services

  This chapter does not explain how to set up and maintain a LAN (Local Area Network), nor does it explain how to install or configure the ARPA services which facilitate cluster communications.

  For those tasks, refer to the following Hewlett-Packard manuals:

  - For an introduction to HP-UX networking services
    - *Networking Overview*
  - For hardware:
    - *HP 98643A LAN/300 Link LANIC Installation Manual* (for a Series 300 server or cluster client)
    - *HP Precision Bus LAN Interface Controller (LANIC) Installation Manual* (for a Model 815 cluster client)
    - *HP CIO LAN Interface Controller (LANIC) Installation and Reference Manual* (for a Series 800 Server)
    - *LAN Cable and Accessories Installation Manual*
  - For software:
    - *Installing and Administering LAN/9000 Series 300*
    - *Installing and Administering LAN/9000 Series 800*
    - *Installing and Administering ARPA Services*

- HP-UX System Administration tasks that are unchanged

  HP-UX clusters are designed to behave as much like standalone systems as possible. There are still significant differences which you as a system administrator need to know about.

  This chapter concentrates on those tasks or aspects of tasks where you must do something different when administering a cluster from what you would do if you were administering a standalone system.

Creating a Cluster

To create a cluster, you need to do the following:

1. Decide what kind of cluster you want.

   See "What Kind of Cluster?".

2. Make sure you have all the hardware and software you need.

   See "Cluster Prerequisites".

3. Make sure the network is properly set up.

   See "Networking Prerequisites".

4. Gather network configuration information.

   See "Filling Out the HP-UX Cluster Information Sheet".

5. Configure the cluster server.

   You'll use the **SAM** (System Administration Manager) utility to do this. See "Configuring the Cluster Server".

6. Add Cluster Clients.

   You'll use the **SAM** (System Administration Manager) utility to do this. See "Adding Cluster Clients".

What Kind of Cluster?

The first thing you need to do is decide what kind of cluster you are going to set up.

As of release 7.0, HP-UX supports two kinds of clusters, **mixed** (or **heterogeneous**) and **homogeneous.**

A **homogeneous cluster** consists of:

- A Series 300 cluster server.
- One or more Series 300 cluster clients.

A **mixed cluster** consists of:

- A Series 800 cluster server.
- One or more Series 300 cluster clients and/or one or more Model 815 cluster clients.

As of Release 7.0, the following restrictions apply:

- A homogeneous cluster must consist only of Series 300 machines.
- A mixed cluster must have a Series 800 machine (not an 815) as cluster server.
- The Model 815 is the only Series 800 machine that can be a cluster client.
- Any cluster including Series 800 machines is treated as a mixed cluster.

  When you use SAM to configure a Series 800 machine as a cluster server, SAM configures the server for both Series 300 and Series 800 cluster clients.

  Thus a "mixed cluster" could consist entirely of Series 800 machines; or of a Series 800 server and Series 300 cluster clients; or of a Series 800 server and both Series 300 and Series 800 cluster clients.

Cluster Prerequisites

To create a cluster you must have:

- For a Series 300 server (homogeneous cluster):
  - Model 319C+, 330, 350, 360, or 370 as the server.
  - At least 8 Mbytes RAM.
  - At least one disk drive of 130 Mbyte or greater capacity.
  - The following products installed:
    - Version 7.0 or later HP-UX (including the AXE and PE products).
    - The LAN/9000 product.
    - The ARPA Services/9000 product.
- For a Series 800 server (mixed cluster):
  - Model 825, 835, 840, 850 or 855 as the server.
  - At least 8 Mbytes of RAM.
  - At least one disk drive, of at least 300 Mbyte capacity.
  - The following products installed:
    - Version 7.0 or later HP-UX for all computers in the cluster (Series 300 and Series 800).

      If your Series 800 server will have Series 300 cluster clients, and is already running an earlier version of HP-UX, you will probably have to *re-install* HP-UX (rather than merely updating to 7.0). See "Configuring the Cluster Server" later in this chapter.
    - LAN/9000.
    - NS/9000.
    - ARPA Services/9000.

- For a Series 300 cluster client:
  - Rev B or later boot ROM
  - LAN hardware
  - At least 3 Mbytes of RAM
- For a Series 800 cluster client:
  - A Series 800 computer as the cluster server
  - A Model 815 computer as the cluster client
  - LAN hardware

Networking Prerequisites

The following guidelines are for those who already understand how to set up and manage a network on HP-UX, using LAN and ARPA services. If you don't know how to do that, don't go on. Stop and read the networking manuals listed at the beginning of this chapter (under "What This Chapter Does Not Cover"). Start with the *Networking Overview,* which is a guide to HP-UX networking products and documentation.

Networking Rules for an HP-UX Cluster

- All cluster nodes must be on a single LAN. This means:
  - Cluster nodes cannot communicate *with each other* via a gateway, (but a cluster can communicate *with other networks* or via a gateway).
  - Cluster nodes *can* communicate with each other over bridges and repeaters.

In addition, we recommend the following:

- Dedicate a LAN to the cluster.
- Use the cluster server, rather than a cluster client, as the gateway to other networks.

  See "Connecting Your Cluster to Another Network" later in this chapter.

Networking Names

Networking services use the following four types of name to identify network nodes (individual machines within the network):

uname The node's SVID/SVVS (System V Interface Definition and System V Validation Suite) name.

Must be 8 characters or less.

hostname The node's ARPA hostname.

Consists of a *name* portion plus *domain* extensions.

nodename The node's NS nodename (if you are using NS services).

cname The node's cluster node name.

To make everything work correctly, **uname,** the *name* portion of **hostname** and **cname** must all be the same. If you are using NS services, **nodename** should preferably be the same as well, but it's not essential.

Because of this, the name you pick must be 8 characters or less, since 8 characters is the maximum allowed by SVID/SVVS, even though the theoretical maximum for each of the other names is greater.

When you configure the cluster server, SAM will verify that these names match as they should.

Filling Out the HP-UX Cluster Information Sheet

Now you need to gather information about the computers that will be in your cluster. Specifically, you need the following for the server and each cluster client:

- ARPA hostname
- Internet Address
- LAN card information:
  - for a Series 300
    - Link Level Address
    - Select Code
  - for a Series 800
    - Station Address
    - Logical unit (lu) number

Link Level Address and **Station Address** refer to the same thing, but the former is more common on the Series 300, and the latter on the Series 800.

The following pages contain blank forms to help you organize this information. The first few lines of the first form are filled out for a sample mixed cluster.

Use the example and the explanations that follow it to help you fill out a form for your cluster.

Note Be sure to read the explanations following the blank forms before you start filling out the form: in some cases the information you need will be on the SAM screen when you go to configure the cluster server or add the client; in others you'll need to do some research. The explanations make clear what's necessary in each case.

Table 10-1. Sample Cluster Information Sheet

Workstation	ARPA Hostname	Internet Address	Cluster LAN card's Link Level Address (Station Address) and Select Code (Logical Unit number)
cluster server	cserver	`192.25.204.1`	0x`080009004a11`/lu=0
client #1	client1	`192.25.204.2`	0x`0800090044ff`/SC=21
client #2	client2	`192.25.204.3`	0x`08000900a63f`/lu=0
client #3	client3	`192.25.204.4`	0x`08000902087a`/SC=21
client #4			
client #5			
client #6			
client #7			
client #8			
client #9			
client #10			

Note

`cserver` is a Model 835 and is the cluster server; `client1` and `client3` are Model 319 clients, and `client2` is a Model 815 client.

Cluster Information Sheet

Workstation	ARPA Hostname	Internet Address	Cluster LAN card's Link Level Address (Station Address) and Select Code (Logical Unit number)
server			/
client #1			/
client #2			/
client #3			/
client #4			/
client #5			/
client #6			/
client #7			/
client #8			/
client #9			/
client #10			/
client #11			/
client #12			/
client #13			/
client #14			/
client #15			/
client #16			/
client #17			/
client #18			/
client #19			/
client #20			/

Cluster Information Sheet

Workstation	ARPA Hostname	Internet Address	Cluster LAN card's Link Level Address (Station Address) and Select Code (Logical Unit number)
client #21			/
client #22			/
client #23			/
client #24			/
client #25			/
client #26			/
client #27			/
client #28			/
client #29			/
client #30			/
client #31			/
client #32			/
client #33			/
client #34			/
client #35			/
client #36			/
client #37			/
client #38			/
client #39			/
client #40			/

ARPA Hostname

Keep this name to 8 characters or less and make it unique within the cluster. See above, under `Networking Names`.

You can leave this blank for the cluster server if the server is already configured into a network: SAM will get the name from the `/etc/hosts` file.

Internet Address

You can leave this blank if it's already recorded in `/etc/hosts` on the server.

Unique network addresses make it possible for a network to communicate with other networks around the world. If your network has not been assigned a unique network address, you can obtain a **Class C** internet address by contacting:

Network Administration Office
Information Networks Division
Hewlett-Packard Company
19420 Homestead Road
Cupertino, California 95014
(408/447-3444)

If you use a network address not assigned by the Network Administration Office and you then need to link with other networks, you may need to change all the addresses in your network.

Internet addresses are usually represented in the form

n.n.n.n

where *n* is a number from 0 to 255, inclusive. (This is referred to as **decimal dot** notation.) For example:

```
192.6.2.9
```

Note This is an example of a **Class C** address. It is likely, but not certain, that your address will be a Class C address. See *Installing and Administering LAN/9000 Series 300,* or *Installing and Administering LAN/9000 Series 800* for an explanation of the three classes of internet address.

In the above Class C example, the high order three numbers (`192.6.2`) are the **network address.** The low order number (`9`) is the **host address.** All nodes in a network share the same network address and each node has a unique host address.

Once you have a network address for the cluster, assign a unique host address to each of the cluster nodes. Do not use leading zeros: a leading zero indicates an octal number, not a decimal number.

In a Class C address, you can assign any number from 1 to 254, inclusive, as the host address.

Caution Do NOT assign 0 or 255 as host addresses in a Class C address. These are reserved addresses. Refer to *Installing and Administering LAN/9000 Series 300* or *Installing and Administering LAN/9000 Series 800* for a full discussion of the rules governing internet addresses.

Link Level Address/Station Address.

The last item on the information sheet is the LAN card's link level address/select code (Series 300) or station address/logical unit number (Series 800).

Fill this in or leave it blank as directed below:

For a Server That Only One LAN card

Leave this blank: SAM will get the number for you.

For a Series 300 Server That Has More Than One LAN Card

You need the link level address and select code of the cluster's LAN card.

You will be able to get a list of the LAN cards on your system, with their select codes and link level addresses, by pressing the `Help` key when you get to the `Link Level Address` field of the `Create an HP-UX Cluster` screen in SAM. If you want to see this information before you get to the SAM screen, you can use the `landiag` utility:

1. Enter

   `landiag`

2. The program displays available commands. Enter

   `l` *(the letter l, for lan)*

3. The program displays another list of commands. Enter

   `d` *(for display)*

4. The program displays information for `Device file`, `Select code`, `current state` and `LAN Interface address, hex`. The `LAN interface address` is the link level address. Write down everything except the leading 0x. For example, if this is what you see,

   ```
   LAN Interface address, hex    = 0x08000903F637
   ```

   then your link level address is 08000903F637

Which LAN card?.

Knowing the link level addresses and select codes may not be enough if don't know which LAN card is the one attaching the server to the cluster.

The following guidelines should help:

- The default select code for a LAN card supplied with the system is 21, so if you have just installed a second LAN card for the cluster, the card you want is probably the one whose select code is *not* 21.
- If you're still not sure, turn off power to your machine, trace the cluster's cable back to the LAN card, and read the link level address off the card.

Write the select code and link level address of the cluster's LAN card on the cluster information sheet.

For a Series 800 Server That Has More Than One LAN card

Write down the station address and logical unit (lu) number of the cluster's LAN card. Follow these steps:

1. Look at the backplane of your computer to get the slot number of the LAN card that is connected to the cluster LAN.
2. Write down the hardware path to the LAN card.

The hardware path for a LAN card is in one of the following forms:

- For an 825, 835 or 840:

  module_number.CIO_slot_number

  module_number is the Mid-bus slot number multiplied by 4.

- For an 850 or 855:

  bus_converter/module_number.CIO_slot_number

  module_number is the Mid-bus slot number multiplied by 4.

For a fuller explanation of hardware addressing, look up your computer in the hardware-specific appendixes to the manual *Installing and Updating HP-UX.*

3. Log in as superuser
4. Change directories:

```
cd /usr/diag/bin
```

5. Run the system diagnostics.

Caution This utility includes destructive diagnostics. Use it *only* as directed below.

Enter

```
/usr/diag/bin/dui
```

6. You'll see a prompt like this:

```
DUI 1>
```

Enter

`run landad pdev=`*path* `section=3`

Where *path* is what you wrote down in step 2.

For example:

```
run landad pdev=8.4 section=3
```

7. Find the entry labelled `NOVRAM (permanent) station address`, for example:

```
NOVRAM (permanent) station address      = $08-00-09-01-91-E5.
```

The station address is everything between the dollar-sign and the period; in this example

```
0800090191E5
```

8. Write down the LAN card's station address and logical unit number on the cluster information sheet.
9. To get out, type

   ```
   exit
   ```

10. Write down the logical unit number.

    Logical unit numbers are sequential, so the first LAN card has a logical unit number of zero, the second LAN card is 1; and so on.

For a Model 815 that currently has a running system on it

You can get the station address and logical unit number (lu) as follows:

1. Get the slot number of the LAN card.
   a. Look at the backplane of the computer. The LAN card occupies two slots, usually 9 and 10.
   b. Multiply the higher slot number by 4. This gives you the module number.

      For example, if the LAN card occupies slots 9 and 10, the module number is 40.
2. Log in as superuser
3. Change directories:

   ```
   cd /usr/diag/bin
   ```

4. Run the system diagnostics.

Caution This utility includes destructive diagnostics. Use it *only* as directed below.

Enter

```
/usr/diag/bin/dui
```

5. You'll see a prompt like this:

   ```
   DUI 1>
   ```

   Enter

   run landad pdev=*path* section=3

   Where *path* is what you wrote down in step 1.

   For example:

   ```
   run landad pdev=40 section=3
   ```

6. Find the entry labelled `NOVRAM (permanent) station address`, for example:

   ```
   NOVRAM (permanent) station address      = $08-00-09-01-91-E5.
   ```

   The station address is everything between the dollar-sign and the period; in this example

   ```
   0800090191E5
   ```

7. Write down the station address on the cluster information sheet.
8. To get out, type

   ```
   exit
   ```

9. Write down the LAN card's logical unit number, which should be 0 (you should have only one LAN card for a cluster client).

For a Model 815 that has no running system on it

In this case you need to get the station address from the LAN card itself.

1. Make sure that the computer and any attached peripherals are powered off.
2. Remove the LAN card, following directions in the *HP Precision Bus LAN Interface Controller (LANIC) Installation Manual.* That manual also tells you how to determine the station address from the LAN card.
3. If you can use the LANIC manual to determine the address, do so. If for some reason you don't have the manual, the following should help.

   The station address is a 12-digit hexadecimal number. The first half of the number is fixed; the second half is what you'll find on the card.

   The first half of the number is

   ```
   080009
   ```

   Now look at the card. The second half of the number is on a label on top of the NOVRAM, which is an IC (integrated circuit) at location U408, roughly in the middle of the board. It consists of six hexadecimal digits, for example

   ```
   4F-E2-54
   ```

   Append these digits to the fixed portion of the number, and write the result on your cluster information sheet, for example

   ```
   0800094FE254
   ```

4. Put the card back, following directions in the LANIC manual.

Configuring the Cluster Server

Configuring a cluster server means converting a standalone system to be a server for a cluster. It involves substantial changes to the kernel and the file system, *and it is irreversible,* in that there is no automated way to "unconfigure" a server once it's been configured. Be quite sure that you want this system to be a cluster server before you embark on the conversion.

Configuring the server involves bringing down the system, so if this is a multi-user system, you will probably want to plan to do the configuration at a time when the fewest possible users will be inconvenienced.

The procedure varies depending whether your cluster will include both Series 300 and Series 800 machines, or machines of only one type (Series 300 or Series 800). The table that follows shows the possible combinations and what to do in each case:

Computers in Cluster	What to do
Series 300 server; series 300 clients	1. Configure the server. See "Using SAM to Configure the Server". 2. Add clients. See "Adding Cluster Clients".
Series 800 server; Model 815 clients *only*	1. Read "Caution" below. 2. Configure the server. See "Using SAM to Configure the Server". 3. Add clients. See "Adding Cluster Clients".
Series 800 server; Series 300 clients	1. Reinstall HP-UX. See "Reinstalling HP-UX (Series 800 Only)". 2. Configure the server. See "Using SAM to Configure the Server". 3. Update the server. See "Updating the Server (Series 800 Only)". 4. Add clients. See "Adding Cluster Clients".

Computers in Cluster	What to do
Series 800 server; Series 300 clients and Model 815 clients	1. Reinstall HP-UX. See "Reinstalling HP-UX (Series 800 Only)". 2. Configure the server. See "Using SAM to Configure the Server". 3. Update the server. See "Updating the Server (Series 800 Only)". 4. Add clients. See "Adding Cluster Clients".

Caution

800 server with 815 clients only: If you are quite sure you will *never* add Series 300 clients to your cluster, then it is sufficient to configure the server and add the Model 815 clients. However, if there's any possibility you may want to add Series 300 clients in future, then you should follow the procedure for "Series 800 server; Series 300 clients and Model 815 clients".

Reinstalling HP-UX (Series 800 Only)

- *If your cluster will contain only Series 300 machines, or only Series 800 machines,* skip to "Using SAM to Configure the Server".
- *If you are configuring a Series 800 machine to be server for Series 300 cluster clients,* then you probably need to reinstall.

Reinstalling will ensure that your root disk is not broken up into small sections (or **partitions**), but has only three sections, of which the largest, covering most of the disk, will contain the / (`root`) file system. This section must be section number 13.

Do You Really Need to Reinstall?

You *do not* need to reinstall if your root disk already uses section 13 for the / (`root`) file system. You probably know whether or not this is the case, but if you're not sure, here's how to tell:

1. Run `mount` to get the name of the device file for the `root` file system.

   Enter

   ```
   mount
   ```

2. You'll see a list of file systems with device file names and other information.

   Find the entry with a / in the first column, for example:

   ```
   / on /dev/dsk/c0d0s0 read/write on Wed Jun 14 02:13:22 1989
   ```

   The device file in this example is `/dev/dsk/c0d0s0`.

   By convention, the number following the `s` at the end of the device file name is the section number, so `c0d0s0` would mean that the / (`root`) file system was on section 0, not section 13, and you'd have to reinstall.

3. To be quite certain of where the root section is, use the `lssf` command:

   Enter

   `lssf` *device_file_name*

   where *device_file_name* is what you got in the previous step.

   For example:

```
lssf /dev/dsk/c0d0s0
```

This will produce a response something like this:

```
disc0 lu 0 unit 0 section 0 address 8.0.0 /dev/dsk/c0d0s0
```

In this example, `section 0` confirms that the `root` file system is on section 0, not section 13, so you'd need to reinstall.

- If your `/` (`root`) file system is on section 13 of your root disk, proceed to "Using SAM to Configure the Server".
- If your `/` (`root`) file system is not on section 13 of your root disk, follow directions under "How to Reinstall".

How to Reinstall

1. Back up your system and re-install HP-UX, using release 7.0 or later.

   Follow directions in chapter 3 of *Installing and Updating HP-UX*. There are specific directions for reinstalling near the beginning of that chapter.

2. When you have re-installed HP-UX and restored your files from the backup, proceed to the next step, "Using SAM to Configure the Server".

Using SAM to Configure the Server

(If you have not used SAM before, you'll find a guide in chapter 1.)

Caution Once you have configured a system as a cluster server, there is no automated way to undo the substantial changes that SAM makes to the kernel and file system (see "What SAM Has Done to Your System" for details of the changes).

---

1. Log in under a superuser id.
2. Get into single-user mode by entering:

   ```
   shutdown
   ```

3. Wait for the system to make the transition into single-user mode (you'll see a shell prompt).
4. Enter `sam`
5. Get to the `Create an HP-UX Cluster` screen:

   ```
   SAM Main Menu
         ↓
   Cluster Configuration
         ↓
   Create an HP-UX Cluster
   ```

   You'll be asked if you are in single-user mode. If you have shut down the system as described above, you are in single-user mode: type `y` and continue.

   If you did not shut down the system, type `n`, back out to the main menu and exit SAM.

   Now the program will check whether this computer meets requirements for a cluster server.

In case of error

Problem	What to do
Message about missing hardware/software	Read Help screen for details.
	Exit SAM.
	Install and configure missing hardware/software.
	Start again.
Message that you're on a cluster client	Read Help screen.
	If you really want this to become a cluster server, you must remove it from the current cluster, add and configure disk(s) for the file system, then configure the system as a cluster server.
Message about missing files	Look in `/tmp/cluster.log` for the names of the missing files and the filesets they belong to.
	Use the `update(1M)` utility to load the filesets (see chapter 3).

You'll notice that this screen allows you to enter information about cluster clients as well as the server. Unless you're already familiar with the process of setting up a cluster, just fill in the information that applies to the server and use the procedure later in this chapter for adding cluster clients.

6. Enter `Node Name` or accept default.

   There will probably be a name displayed in this field: SAM gets it from the `/etc/hosts` file. Don't change it.

   If you see a window prompting you for a hostname, type the cluster server's ARPA hostname from your cluster information sheet.

   The node name can be 1 to 8 ASCII characters; must not be `default`, `localroo`, `remotero` or begin with `HP`; and must not contain spaces, newline characters, or pound signs (`#`).

   If you have two LAN cards, and one is configured to communicate with systems outside the cluster, you use its ARPA hostname here. See "Connecting Your Cluster to Another Network", later in this chapter, for a full explanation.

7. Enter `Internet Address` or accept default.

   There's probably an address already displayed in this field: it is the internet address associated with the cluster node name shown on the screen, derived from the `/etc/hosts` file. You can't change it.

   If there's nothing displayed here, type the cluster server's internet address from your cluster information sheet.

8. Enter `Link level address` (station address) or accept default.

   This field is always filled in, but if your cluster server has more than one LAN card, the number displayed may not be the one you need.

   - If you have only one LAN card, the number already filled in is correct. *Do not change it.* Write the number onto your cluster information sheet.
   - If you have more than one LAN card, check the number on the screen carefully against what you have on your cluster information sheet. If they don't match, type the correct value from your sheet into the field on the screen.

     To check the link level addresses and select codes (or station addresses and logical unit numbers) of the LAN cards on your system, press `Help`.

   If you don't know which is the right address, and you did not write it on your cluster information sheet, you need to go back to the section "Link Level Address/Station Address" earlier in this chapter.

9. Accept default `Machine`.

   This defaults to the correct value: `s800` or `s800`, depending on which type of computer the server is.

10. Enter `# of CSPs` or accept default.

    This is the number of cluster server processes (CSPs) you wish to run.

    In general, the more cluster nodes in your cluster, the greater this value should be, but it's best to take the default (4) for now until you've had a chance to monitor performance and tune the system. The "Cluster Concepts" chapter of the *HP-UX System Administration Concepts* explains CSPs and how they work.

11. Press `Perform Task`.
12. You'll see a warning that it's difficult to convert the cluster server back to a standalone machine once the cluster configuration has been done.

    Assuming you do want to configure this computer as a cluster server, respond `y` to continue.
13. If your server is a Model 850 or 855, you'll be prompted for the name of your `uxgen` input file. By default, this file is `/etc/conf/gen/S800`, but if you have changed the name, respond with the current name here.
14. Now you'll be asked if you want SAM to reboot the system. Respond `y`: the changes will not take effect until you have rebooted the system.

Now SAM performs the tasks necessary to convert your system to a cluster server. It creates CDFs, builds a cluster kernel, creates and/or updates the necessary ARPA files, and creates the file `/etc/clusterconf` with an entry for `localroot` (the cluster server). You'll see messages telling you what's happening. See "What SAM Has Done to Your System" for more details.

The customization will take a few minutes, after which the system automatically reboots (unless you replied `n` to the question about rebooting).

Your machine is now configured as a server to support an HP-UX cluster. Before you proceed to add clients, you should:

1. Read the next section, "What SAM Has Done to Your System".
2. Print and save a hard copy of the file `/tmp/cluster.log`, which contains a complete list of all the **context-dependent files (CDFs)** which SAM has created on your system.

   A context-dependent file is a file whose contents differ depending on which member of the cluster is using it: see the Cluster Concepts chapter in the *System Adminisration Concepts* manual for a full explanation.

What you do next depends on the make-up of your cluster. The table that follows shows the possible combinations and what to do in each case.

Computers in Cluster	What to do
Series 300 server; series 300 clients	Add clients. See "Adding Cluster Clients".
Series 800 server; Model 815 clients *only*	1. Read "Caution" below. 2. Add clients. See "Adding Cluster Clients".
Series 800 server; Series 300 clients	1. Update the server. See "Updating the Server (Series 800 Only)". 2. Add clients. See "Adding Cluster Clients".
Series 800 server; Series 300 clients and Model 815 clients	1. Update the server. See "Updating the Server (Series 800 Only)". 2. Add clients. See "Adding Cluster Clients".

Caution **800 server with 815 clients only:** If you are quite sure you will *never* add Series 300 clients to your cluster, then it is sufficient to configure the server and add the Model 815 clients. However, if there's any possiblity you may want to add Series 300 clients in future, then you should follow the procedure for "Series 800 server; Series 300 clients and Model 815 clients".

What SAM Has Done to Your System

You have now configured your cluster server. SAM has done the following:

- Created the `/etc/clusterconf` file

  There will be two lines in the `/etc/clusterconf` file. The first line contains your cluster server's link level address. The second line has some cluster server node information.

  See `clusterconf(4)` in the *HP-UX Reference* for more information.

- Turned certain files into CDFs

  For a list of the files that are now CDFs, refer to the file `/tmp/cluster.log`. If you have not already printed this file, do so now and save it for future reference.

- Modified the kernel to include cluster configuration.

  `/hp-ux` is now a CDF. Your cluster server's version of the kernel resides in the file `/hp-ux+/`*server_nodename*. For example, if your server's cluster node name is `cserver`, the kernel will be in `/hp-ux+/cserver`.

  This kernel has the drivers and the parameter values needed by a cluster server.

  The old kernel has been saved as `/SYSBCKUP` and as `/hp-ux+/standalone`.

  - On a Series 300:

    The file `/etc/conf/dfile+/`*server_nodename* matches the kernel.

  - On a Series 800:

    The file `/etc/conf/gen/S800+/`*server_nodename* matches the kernel.

  The kernel is built from the existing compiled kernel (`/hp-ux`), except in the case of an 850/55, when the kernel is built from the `uxgen` input file: `/etc/conf/gen/S800` or whatever name you gave SAM when you were prompted on the `Create an HP-UX Cluster` screen.

  The old S800 file is saved as `/etc/conf/gen/S800.BCKUP`

- Put an entry for the cluster server in each of the following files (unless the entry was already there):
  - /etc/hosts
  - /etc/hosts.equiv
  - $HOME/.rhosts (root's home directory)
  - /etc/X0.hosts (if it exists)
- Modified the line RBOOTD_DEVICES= in /etc/rc to point to the correct device file for the LAN card.

Updating the Server (Series 800 Only).

If your server is a Series 300, or a Series 800 machine that will serve ONLY Series 800 cluster clients, skip this step and proceed to add cluster clients.

If your Series 800 server will have Series 300 cluster clients, you must now install the Series 300 operating system on your Series 800 server.

Caution You *must* convert a Series 800 computer to a cluster server before you install the Series 300 operating system on it.

---

1. Turn to chapter 4 of *Installing and Updating HP-UX*.
2. Update your system, using the *Series 300* update tapes.

   Choose the option to `Load Everything from Source Media`.

   So long as you have already configured your machine as a cluster server, the software specific to the 300 will co-exist with the Series 800 software without over-writing it.

After you have installed the Series 300 operating system, you are ready to add cluster clients. Continue with "Adding Cluster Clients" below.

Adding Cluster Clients

Use the SAM utility to add a cluster client. *Do this on the CLUSTER SERVER.*

Prerequisites

Before adding cluster clients make sure the following conditions exist:

- The cluster server is configured.

  See "Configuring the Cluster Server" earlier in this chapter.
- Each computer to be made a cluster client is a Model 815 or a Series 300 machine.
- If any cluster client is an 815, the server must be a Series 800 machine.
- Each Series 300 cluster client has Rev B or later boot ROM.
- Each cluster client has at least 3 Mbytes of RAM.

Note Although the following pages provide separate procedures for adding a Series 300 and a Model 815 client, you can add clients of both types at the same time: the SAM screen allows you to add as many clients as you need to in a single transaction.

Before you can start adding clients, you need to run SAM *on the cluster server,* and get to the `Add Cluster Clients` screen in SAM:

```
SAM Main Menu
     ↓
Cluster Configuration
     ↓
Add Cluster Clients
```

In case of error

Problem	What to do
Message that this is not a member of a cluster	You have not configured this computer as a cluster server. Check that you are on the cluster server. If you have not yet configured a cluster server, do so now. See "Configuring the Cluster Server" earlier in this chapter.

Adding a Series 300 Cluster Client

Summary of steps

1. Get the link level address *(on client)*
2. Use SAM to add the cluster client *(on server)*
3. Boot the client *(on client, after all clients have been added)*

Getting the Link Level Address (Series 300)

Do this step on the CLUSTER CLIENT.

You will need your cluster information sheet.

1. If the computer to be made a cluster client has a running system on it, go to the console and type

   ```
   shutdown -h 0
   ```

   then wait for the "halted" message.

Note This applies to the *cluster client,* not to the server. The server should be up and running.

---

2. Turn off power to the computer.
3. If any disk attached to the *cluster client computer* has a bootable system on it, it's safest to remove the disk, or at least turn it off.

   If you need to leave the disk attached (if this client is going to double as a standalone system), you will need to be careful to pick the right operating system when you go to boot the client over the LAN: it will appear under the `LAN` statement. "Booting the Client (S300)", later in this chapter, gives details.
4. Turn on power to the computer and go immediately to the console.
5. Press the space bar and continue to hold it down until you see the word `keyboard` on the left of the console screen.

   You will see a screen similar to the one shown in Figure 10-1.

```
Copyright 1987,
Hewlett-Packard Company.
All Rights Reserved.

BOOTROM Rev. C
Bit Mapped Display
MC68050 Processor
Keyboard
HP-IB
HP98620B
HP98644 at 9
HP98625 at 14
HP98643 at 21, 080009000001
4182016 Bytes







SEARCHING FOR A SYSTEM (RETURN To Pause)
RESET To Power-Up
```

Figure 10-1. Sample Boot ROM Display

A LAN card will show up in the format:

lancard_partnum `at` *select_code*, *link-level-address*

(for example, `HP98643 at 21, 080009000001`).

6. Write down the cluster LAN card's link level address on the cluster information sheet.
7. Leave the cluster client in this state. (This is called **attended boot mode.**) You will come back later to finish the boot sequence.

Using SAM to Add the Cluster Client (Series 300)

Do this on the CLUSTER SERVER.

Prerequisites.

1. If the cluster server is a Series 800 computer, you *MUST* update it with the *Series 300* operating system before you can add Series 300 clients.

   See "Updating the Server (Series 800 Only)" earlier in this chapter, and go back to "Configuring the Cluster Server" if you need to check the order of events.

2. You should be in SAM on the server, on the `Add Cluster Clients` screen. If you are not, here's a map of how to get there:

```
        SAM Main Menu
              ↓
    Cluster Configuration
              ↓
     Add Cluster Clients
```

Procedure.

1. Type the `Client Name`. This is the ARPA hostname for this cluster client from your cluster information sheet.

   The name must be at least 1 and up to 8 characters (any ASCII characters) and must be unique to the network.

Note If this node will be a gateway to another network that already exists, this name must be the same as the existing ARPA hostname and HP-UX system hostname. This configuration is not recommended. See "Networking Prerequisites" earlier in this chapter.

---

2. Enter the cluster client's ARPA internet address.

   This is the internet address for this cluster client from your cluster information sheet.

   If an address is displayed here, it is the address associated with the ARPA hostname you entered under `Client Name` (see step 1 above) and you can't change it.

3. Type the cluster client's LAN card link level address.

   You wrote this number down on your cluster information sheet. (If you didn't, and you don't know the link level address, go back to "Getting the Link Level Address (Series 300)" now.)

   Do not type the select code, just the link level address.

4. Enter the machine type: `s800`.
5. Add the information for the next cluster client.

   (If the next cluster client to be added is a Model 815, use the procedure under "Adding a Model 815 Cluster Client" later in this chapter.)

6. If this is the last cluster client you need to add, do the following:
   a. Press `Perform Task`.
   b. Confirm that you want to complete this transaction by typing `y` when prompted.
   c. Wait for the message `The cluster clients were added as requested.`

Booting the Cluster Client (Series 300).

Do this after you have entered all your cluster clients onto the SAM `Add Cluster Clients` screen.

1. *Go back to the cluster client,* which you left in attended boot mode.

   The console should have a set of entries on the right side looking something like this:

```
LAN, 21, cserver
 1H SYSHPUX
 1D SYSDEBUG
 1B SYSBCKUP
```

2. Choose the `SYSHPUX` option that appears under the `LAN, ...` statement.

In case of error

Problem	What to do
Boot fails	Check `/tmp/cluster.log` to see if any errors occurred while you were using SAM to add the client. Check that `rbootd` is running on the server. Reboot (turn power off and on).

3. When the cluster client has booted, you will see a login prompt. Log in. Remember that you need to use the cluster server login and password.

Are you booted to the wrong system?.

If you did not leave the cluster client in attended boot mode (as described earlier in this chapter, under "Getting the Link Level Address (S300))", and you did not remove or turn off any disk drive containing a bootable system, then you are probably booted to the system on the disk.

Log in. If this is an HP-UX system, type:

```
getcontext
```

- If the returned string contains the word `remoteroot`, you are booted to the right system.
- If the response is

  ```
  not found
  ```

  or if the returned string contains the word `standalone` or `localroot`, then you are booted to the wrong system. If you are booted to the wrong HP-UX system, do the following:

1. Enter

   ```
   shutdown -h 0
   ```

2. When you receive the "halted" message, turn the cluster client's power off and on and boot it in attended mode.
3. If you do not see the new operating system entry, turn power off and on again to restart the boot ROM and see if it now sees the cluster server.
4. If you still do not see the cluster server entry, check that `rbootd` is running on the server.
5. If `rbootd` is running, you may have entered an incorrect link level address: start again with "Getting the Link Level Address (Series 300)" and verify all the information you have entered.
6. If any of the information you've entered is wrong, you'll need to delete the client and add it back correctly. See "Removing Cluster Clients" later in this chapter.

Adding a Model 815 Cluster Client

Summary of steps

1. Prepare the cluster client *(on client)*
2. Use SAM to add the cluster client *(on server)*
3. Boot the cluster client *(on client, after all clients have been added)*

Preparing the Cluster Client (Model 815)

Do this on the cluster client.

The first thing you need to do is check whether or not you have a standard hardware configuration. You need to know this because SAM is going to build a kernel for your 815 client *on the server*. A template file matching the standard configuration is shipped with every Series 800 that is supported as a cluster server: adding a Model 815 client is simpler if its hardware configuration matches the standard template.

If you bought your 815 as a cluster client and have never used it for anything else, you probably have the standard configuration.

Figure 10-2 is a diagram of the backplane of an 815 in the standard configuration. Go to the back of your 815 and check the placement of the cards.

```
 ----------------------------------------------------------
|     (Mux 5 if installed)  |  (Mux 6 if installed)        |
|                (slot 13)  |  (slot 14)                   |
|----------------------------------------------------------|
|        (PSI if installed) |  (Mux 4 if installed)        |
|                 (slot 11) |  (slot 12)                   |
|----------------------------------------------------------|
|             <------------LAN------------>                |
|                (slot 9)   |  (slot 10)                   |
|----------------------------------------------------------|
|           Mux 2           |                Mux 3         |
|                (slot 7)   |  (slot 8)                    |
|----------------------------------------------------------|
|           HP-IB           |                Mux 1         |
|                (slot 5)   |  (slot 6)                    |
|----------------------------------------------------------|
|    BBU (battery backup)   |                8M memory     |
|                (slot 3)   |  (slot 4)                    |
|----------------------------------------------------------|
|                                                          |
|   computer and 2 on-board RS232 slots A and B.  Console in A.  |
|                 (pseudo slots 0-2)                       |
 ----------------------------------------------------------
```

Figure 10-2. Model 815 Standard Hardware Configuration

Note To match the standard kernel template, cards must be in the slots shown *if they're present.* If you don't have a MUX card at slot 12, for example, you still have a standard configuration so long as slot 12 is empty or contains a memory card.

- If your hardware configuration matches the diagram, you don't need to do anything.

  A kernel source file corresponding to this configuration, `/etc/conf/gen/M815.dskless`, is already in place on the cluster server.

- If your hardware configuration does not match this diagram, you'll need to create a kernel source file matching the 815's actual configuration.

  This file must be *on the server* before you add the 815 as a cluster client.

  If you have a non-standard configuration it's probably because you have been running the 815 as a standalone system and are now going to convert it to a cluster client. In other words, you very likely have a file that matches your current hardware configuration and do not have to create it from scratch.

  In this case you will need to copy the S800 file (usually `/etc/conf/gen/S800`) from the 815 system onto the server. Make sure you don't overwrite the server's S800 file: copy the client's file to a different name, for example `/etc/conf/gen/S800.client3`. The file *must* reside in `/etc/conf/gen`.

Now get the 815 ready to be booted as a client.

1. Shut down and turn off the computer.
   - If the 815 has a system running on it, do the following *on the 815 client:*
     a. Shut down the 815.

        Make sure all users are off the system, then go to the console and type:

        ```
        shutdown -h 0
        ```

        then wait for a message like this:

        ```
        HALTING IN TIGHT LOOP--OK TO PRESS RESET
        ```

     b. Turn off power to the computer.
   - If the Model 815 has no operating system, simply make sure that the computer is turned off.

2. If any disk attached to the *Model 815 computer* has a bootable system on it, remove the disk, or at least turn it off.

   This is not an absolute requirement, but it will make things much simpler.

3. Leave the cluster client powered off. You will come back later to boot it over the LAN.

Using SAM to Add the Cluster Client (Model 815)

Do this on the CLUSTER SERVER.

Prerequisites.

You should be in SAM on the server, on the `Add Cluster Clients` screen. If you are not, here's a map of how to get there:

```
SAM Main Menu
      ↓
Cluster Configuration
      ↓
Add Cluster Clients
```

Procedure.

1. Type the `Client Name`. This is the ARPA hostname for this cluster client from your cluster information sheet.

   The name must be between 1 and 8 characters (any ASCII characters) and must be unique to the network.

Note If this client will be a gateway to another network that already exists, this name must be the same as the existing ARPA hostname and HP-UX system hostname. This configuration is not recommended. See "Networking Prerequisites" earlier in this chapter.

---

2. Enter the cluster client's ARPA `Internet Address`.

   This is the internet address for this cluster client from your cluster information sheet.

   If an address is displayed here, it is the address associated with the ARPA hostname you entered under `Client Name` (see step 1 above) and you can't change it.

3. Type the cluster client's LAN card `Link Level Address` (station address).

   You wrote this number down on your cluster information sheet. (If you didn't, and you don't know the internet address, go back to "Filling Out the Cluster Information Sheet" now.)

4. Enter the system type: `s800`.
5. Now you'll be prompted for the name of a template file from which SAM can build the client's kernel.

   The file `/etc/conf/gen/M815.dskless` corresponds to the standard configuration. You checked your 815's configuration earlier—if you didn't, and you're not sure whether or not you have the standard configuration, go back to "Preparing the Cluster Client (Model 815)" now.

   If your 815 does not have the standard kernel configuration, you should have copied onto the server a file that reflects the actual configuration (see "Preparing the Cluster Client (Model 815)" above): respond to the prompt with the name of that file (for example `S800.client3`). Remember that the file must reside on the server, in the directory `/etc/conf/gen`.
6. Add the information for the next cluster client.

   (If the next cluster client to be added is a Series 300, use the procedure under "Adding a Series 300 Cluster Client" earlier in this chapter.)
7. If this is the last cluster client you need to add, do the following:
   a. Press `Perform Task`.
   b. Confirm that you want to complete this transaction by typing `y` when prompted.
   c. Wait for the message `The cluster clients were added as requested`.

Booting the Cluster Client (Model 815).

Do this after you have added all your cluster clients using the SAM `Add Cluster Clients` screen.

You left the cluster client turned off. To boot it, do the following:

1. Read through this procedure to the end before you do anything, so that you'll know what to expect.
2. Turn power on to the computer.
3. Now one of two things will happen:
   - If you see a prompt like this,

     ```
     IPL>
     ```

     enter

     ```
     primpath n.0.0.0.255
     ```

     where *n* is arrived at by taking the higher number of the two slots the LAN card occupies, and multiplying this slot number by 4.

     For example, if your LAN card is in slots 9 and 10, respond

     ```
     primpath 40.0.0.0.255
     ```

     Go to step 4.
   - If you see a message telling you to press a key on the keyboard, press any key immediately.
     a. The system will ask if you want to boot from the primary path: respond **n**.
     b. Now the system will ask if you want to boot from the alternate boot path: respond **n**.
     c. The system will prompt

        ```
        Enter boot path, command or ?>
        ```

        Respond as follows:

        ```
        p primary n.0.0.0.255
        ```

        where *n* is the number arrived at by multiplying the LAN's even-numbered slot by 4.

        For example, if your LAN card is in slots 9 and 10, respond

        ```
        p primary 40.0.0.0.255
        ```

You are telling the computer to set the primary boot path to the address of the LAN card (and, if the server should go down, to wait for it to come back up, and then reboot).

d. Now you'll see this prompt:

```
PDC>
```

Type the letter `b` (for "boot").

e. The system will ask if you want to interact with IPL.

Respond `y`.

You'll see boot messages and then

```
IPL>
```

In case of error

Problem	What to do
Failed to press a key to interrupt boot *or* accidentally responded `y` to boot from primary or alternate path *or* accidentally responded `n` to `Interact with IPL?`	Turn power off and on and start again.

4. At the `IPL>` prompt, type

```
autoboot on
```

5. You have set up the cluster client to boot automatically over the LAN from the kernel that was built when you configured the client into the cluster (see "Using SAM to Add the Client (Model 815)" earlier in this chapter).

   The client will boot automatically over the LAN every time from now on, until you change the autoboot or primary path setting, or change the hardware configuration.

Now turn power off and on and boot the node using the directions that follow.

6. Turn power on.
7. Allow the boot to complete uninterrupted.

   Do not press a key when you see the message telling you you have 10 seconds to interrupt.

   (If you press a key by accident, just respond `y` when the system asks if you want to boot from the primary path and `n` to `Interact with IPL?`.)

   When you see the login prompt, the cluster client is ready to use.

After Adding Cluster Clients

You can now add terminals and other peripherals for the new cluster clients if you wish.

Because the `/dev` directory is a CDF, any peripherals you attach to a cluster client can be used only by that client, and you must add them while logged into the cluster client. On the other hand, printers, plotters and disks attached to the server can be shared. See "Peripherals: Points to bear in Mind", later in this chapter, for more information.

If you want a cluster client to swap to its own local disk, follow the instructions later in this chapter under "Adding a Local Swap Device". In most clusters the clients will swap to the server's disk space.

Connecting Your Cluster to Another Network

Note Before you do anything described in this section, you must first create your cluster: follow directions earlier in this chapter for configuring the cluster server and adding cluster clients. If you are unfamiliar with networking concepts, you should also study the networking documentation listed near the beginning of this chapter.

Recommendations

A cluster is itself a network. However, you may want the cluster to communicate with other machines which are not part of the cluster.

This means that one of the machines in the cluster will be a LAN **gateway:** that is, it will have two (or more) LAN cards, one connecting it to the cluster and the other(s) connecting it to the other network(s).

If you need to connect your cluster to another network, Hewlett-Packard recommends that you use the cluster server as the gateway, not one of the cluster clients.

Configure the networks as follows:

- Put all the cluster nodes (the cluster server and all the cluster clients) on one LAN.
- Configure the cluster server as the LAN gateway.
- Make the **official host name** in the `/etc/hosts` file the same for both LAN cards.
- Supply a unique **alias** in `/etc/hosts` for each card.

A Worked Example

Your configuration should look like something like this:

```
 ----------        ----------       ----------
| 192.1.1.2|      | 192.1.1.3|     | 192.1.1.4|
|          |      |          |     |          |
| "client1"|      | "client2"|     | "client3"|
 ----------        ----------       ----------
      |                 |                |
      |                 |                |
-------Cluster LAN (192.1.1 Network)-----------
                        |
      LAN card at Select Code 21 (or lower lu)
                        |
                   ----------
                  | 192.1.1.1|
                  | "cserver"| Cluster Server/
                  |          | Gateway Node
                  | 192.2.2.1|
                   ----------
                        |
      LAN card at Select Code 22 (or higher lu)
                        |
-------Non-Cluster LAN (192.2.2 Network)-------
```

This picture shows two networks (network addresses 192.1.1. and 192.2.2).

The cluster network includes three cluster clients (`client1`, `client2` and `client3`) and the cluster server, `cserver`.

`Cserver` is a gateway node, connecting the cluster network (192.1.1) to the other (192.2.2) network.

The LAN card connecting the cluster server to the cluster network is at a lower select code (or logical unit number in the case of a Series 800 computer) than the LAN card connecting to the other network. (This is preferable, but not required.)

To make the two networks co-operate efficiently, you need to do the following:

1. Create the cluster.

   You should have done this already. If you haven't, go back to "Creating a Cluster" earlier in this chapter.

   Do not do either of the following steps until after you have created the cluster.

2. Modify `/etc/hosts`.

   The networks represented on the previous page should be described in `/etc/hosts` as follows:

(Internet Address)	(Official Hostname)	(Alias)	*Headings are not in file*
192.1.1.1	cserver	dserver	
192.1.1.2	client1		
192.1.1.3	client2		
192.1.1.4	client3		
192.2.2.1	cserver	dgateway	

3. Modify /etc/netlinkrc.

Note Do not change /etc/netlinkrc until after you have created your cluster.

To reflect the configuration shown on the previous page, you would modify /etc/netlinkrc as follows:

a. Find the statement

```
case $NODENAME in
    *) ifconfig lan0 'hostname' up
    ;;
esac
```

b. Insert text so that the statement in step 1 now reads:

```
case $NODENAME in
     $ROOTSERVER) ifconfig lan0 dserver up;
                  ifconfig lan1 dgateway up;
    *) ifconfig lan0 'hostname' up
    ;;
esac
```

c. Find the statement

```
case $NODENAME in
     *) # comment
        ;;
```

d. Insert text so that the statement in step 3 now reads:

```
case $NODENAME in
     *) # comment
     $ROOTSERVER);;
     *) /etc/route add default dserver 1
        ;;
```

Removing Cluster Clients

This section explains how to remove cluster clients. (You cannot remove the cluster server).

Use the procedure that follows to do either of the following:

- Remove one or more clients.
- Rename one or more clients.

  To rename a cluster client, you should remove it, using the procedure that follows, then add it back, using directions under "Adding Cluster Clients" earlier in this chapter. Before you go on, read the next section, "Renaming Cluster Clients".

Do this on the SERVER.

1. Log in to the server under a superuser id.
2. Get to the `Remove Cluster Nodes` screen in SAM:

   ```
   SAM Main Menu
        ↓
   Cluster Configuration
        ↓
   Remove Cluster Clients
   ```

3. Select the clients you want to delete by typing `x` next to their names in the `Remove Client` column.
4. Answer `y` or `n` to the question `Remove Client-Specific Files?`.
   - If you reply `y`, the SAM program will check the file system for all CDFs that contain a file with the cluster client's name, and will remove those files.
   - If you answer no, these client-specific files will stay in the cluster server's file system.

   Answer `y` unless you have a specific reason to leave the client-specific files on your system.
5. Press `Perform Task`.

If you have opted to remove the client-specific files, you'll see a message that a background script is running to remove client-specific files. You can leave this screen, and exit SAM altogether if you wish.

If, after pressing `Perform Task`, you decide to remove more clients, you must wait until the script has finished. The script will write a message at the end of `/tmp/cluster.log` when it's finished.

6. Exit SAM.

Renaming Cluster Clients

The easiest way to rename a cluster client is to remove it and then add it again with a different name. Follow these steps:

1. If you have customized any of the client's elements in the cluster's CDFs, you may want to save them by copying them to a different name.

   Let's say you are going to rename `client1` to `george`, and that you have customized `client1's` `/etc/inittab`.

   To save the customized copy of `/etc/inittab`, you might decide to copy it to `/tmp/inittab.george`:

   ```
   mv /etc/inittab+/client1 /tmp/inittab.george
   ```

   Since this command makes explicit reference to `client1's` element of `/etc/inittab`, you can issue it on the cluster server or on `client1` with the same effect.

2. Use the SAM program on the cluster server to remove the cluster client and all its elements in the cluster's CDFs.

   Follow the directions under "Removing a Cluster Client" in the previous section. Answer `y` to the question `Remove Client-Specific Files?`.

3. Edit the `/etc/hosts` file. Change the name field in the cluster client's entry in `/etc/hosts` to the new name.

4. Using the SAM program, add the cluster client under the new name. Follow the directions under "Adding Cluster Clients" earlier in this chapter.

5. If you saved any files in step 1, copy those files back to the CDFs where they belong.

   For example, if you had saved the client's version of `etc/inittab` as `/tmp/inittab.george`, and then renamed the client to `george`, you would issue the command:

   ```
   mv /tmp/inittab.george  /etc/inittab+/george
   ```

Administering an HP-UX Cluster

Once your cluster is up and running, administering it is not very different from administering a standalone machine. Tasks that differ significantly from the standalone case, or are unique to cluster administration, are covered later in this chapter.

However, one question that is likely to keep coming up in your mind is *where* to perform a given task—on the cluster server or on a client?

The tables that follow should help.

Table 10-2. Routine Tasks: When and Where to Perform Them

What	When	Where
Configure cluster	Set-up	Server (1)
Add clients	Cluster set-up, expansion	Server (2)
Remove clients	Cluster restructuring	Server (3)
Boot server	Set up, maintenance	Server
Boot client	Set up, maintenance	Client
Set time and date	Initial boot	Any cnode (4)
Create first full archival backup	After first boot	Server (5)
Do incremental backup	Daily	Server (5)
Do full backup	Weekly	Server (5)
Create recovery System (S300 only)	After first boot	Server
Create new user accounts	After first boot, expansion	Any cnode (4)
Check disk usage	Weekly	Server

Notes

1. See "Creating an HP-UX Cluster" earlier in this chapter.
2. See "Adding Cluster Clients" earlier in this chapter.
3. See "Removing Cluster Clients" earlier in this chapter.
4. You can do this on any cluster node (the server or any client) and the result will be global to the cluster.
5. To back up a cluster's files you need to use special options to locate the contents of context-dependent files (CDFs). For example, use the `-hidden` option of the `find(1)` command when backing up with `tcio(1)/cpio(1)`.

   The *HP-UX Reference* has full details of these commands, and there are examples of backing up a cluster in chapter 11 of this manual.

   For performance reasons, you should do backups on the cluster server.

Table 10-3. Event-Driven Tasks

Event	What To Do	Where
Power going out soon	Shut down cluster. Turn power off for server and clients.	Server, clients, peripherals (1)
Bringing cluster back up	Boot server, then clients	Server, clients (2)
Powerfail on server	Reboot clients	Clients (3)
Powerfail on client	Nothing. Client will recover	N/A (3)
Server maintenance needed	Shut down the cluster, power down server	Server (1)
Client maintenance needed	Halt and power down client	Client (4)
Need to send message to all cluster users	Use `cwall(1m)`	Any cnode (5)
Need to send message to all users of one cnode	Use `wall(1m)`	Affected cnode
Files accidentally deleted	Recover files from backup	Server
File system corrupted	Use `fsck (1M)`, recovery system, archives as appropriate	Server (6)
New user arrives	Use SAM to create account	Any cnode (5)
User leaves	Use SAM to delete account	Any cnode (5)
Need to create/change group	Use SAM	Any cnode (5)
System clock wrong	Set clock ahead/back	Any cnode (5)
Server panics	Diagnose problem using the *Troubleshooting HP-UX Systems* manual. Reboot server. Recover files. Reboot clients.	Server, clients

Event-Driven Tasks (continued)

Event	What To Do	Where
Adding a peripheral	Follow directions in *Installing Peripherals* (Series 300) or *Confguring HP-UX for Peripherals* (Series 800). Use SAM if possible.	Cnode peripheral is connected to
Changing kernel configuration	Rebuild/reinstall kernel	Cnode that will use kernel (7)

Notes

1. See "Shutting Down a Cluster" later in this chapter.
2. See "Booting a Cluster" later in this chapter.
3. When a powerfail occurs on a cluster server, all the clients will panic.

   To bring the cluster clients back up after a powerfail on the server, wait for the server to recover (which it should do automatically), then reboot the clients.

   A cluster client that experiences a local powerfail (i.e., one that does not affect the server) will recover automatically without affecting the rest of the cluster.

   To make it less likely that cluster clients will panic during a powerfail on the server, you can increase the values of the clients' kernel parameters `check_alive_period` and `retry_alive_period`, but be aware of the trade-offs:

   - If the cluster server crashes, the clients will have to wait longer before they realize the server is no longer there.
   - If a client crashes, it could lock the entire cluster until the `check_alive_period` and `retry_alive_period` have expired.

   If you decide to change these parameters, you will need to reconfigure the kernel of each affected client. Use SAM to do this. See "Configuring the

Kernel for a Cluster Node" later in this chapter, and chapter 10 of this manual.

4. See "Shutting Down a Cluster Client" later in this chapter.
5. You can do this on any cluster node (the server or any client) and the result will be global to the cluster.
6. To run `fsck(1M)` on the root file system, you must be logged in to the server *directly,* not via **`rlogin`**.
7. See "Reconfiguring the Kernel for a Cluster Node" later in this chapter, and chapter 10 of this manual.

Table 10-4. Miscellaneous Tasks: Where to Perform Them

What	Where
Mount/unmount file systems on HFS	Server
NFS mount/umount	Server, clients (1)
Create file system	Server
Update HP-UX	Server
Install or update applications	Server (2)
Remove filesets	Server
Configure LP spooler	Server (3)
Set up, use UUCP	Server (3)
Modify sytem files	Any cnode (1, 4)

Notes

1. You can do this from any cluster node, and the result will be global to the cluster.
2. Before installing applications in a mixed cluster, see "Installing and Removing Programs in a Mixed Cluster" later in this chapter.
3. Spooling and UUCP are supported only on the cluster server.

   Spooling:

   - Spooled printers can be attached only to the server.
   - Unspooled printers can be attached to any cluster node.
   - Print requests can be spooled from any cluster node, but the scheduler runs only on the server.

   UUCP:

   - All UUCP connections must be on the cluster server.
   - UUCP transfers can be initiated from any cluster node.
4. Make sure that system files are modified by only one user at a time.

A Key Concept: Context-Dependent Files

In order for different machines to use the same file system, a given file may need to have different contents depending on which particular member of the cluster is looking at it. This type of file is known as a **context dependent file** or **CDF.**

As you configure a cluster, HP-UX utilities automatically build the CDFs the system will need, but you may also need to create your own CDFs as you bring new applications onto the cluster.

Instructions for installing a binary file (such as a program file), whose contents must be different depending on what type of computer is using it, are in the section "Installing Programs in a Mixed Cluster". CDFs are covered in detail in the Cluster Concepts chapter of the *HP-UX System Administration Concepts* manual.

Using HP-UX Commands with Clusters

This section lists commands and options of commands that are useful in administering an HP-UX cluster, as well as commands that you *cannot* use on a cluster client (though they work normally on the server). At the end of the section are some notes about things that work differently in a cluster from the way you might expect.

The *HP-UX Reference* and the "Cluster Concepts" chapter of the *HP-UX System Administration Concepts* manual provide more details.

(The numbers in parentheses immediately following the command names refer to the relevant sections of the *HP-UX Reference*.)

Commands Specific to Clusters

showcdf(1)

Expands the path name of any CDF to show the actual name of the file on the current node. Use the -c option to avoid cluttering up the display with the names of files that are not CDFs.

For example, on system cserver, the command

```
showcdf -c /etc/checklist
```

would produce this output:

```
/etc/checklist+/cserver
```

The + shows that etc/checklist is really a directory (formally referred to as a **hidden directory**) containing the cnode-specific file cserver (i.e., this cluster node's version of /etc/checklist).

getcontext(1)

Shows the **context** of the current cluster node.

The context is an ASCII string that identifies this particular cluster node. You need to know the context when creating context-dependent files (see "Installing and Removing Programs in a Mixed Cluster" later in this chapter, as well as the *HP-UX System Administration Concepts* manual).

cwall(1M)

Cluster version of wall(1m).

Use cwall when you want to broadcast a message to all users of all cluster nodes, for example to warn them that the cluster is about to go down. You can use cwall from any cluster node, and everyone using the cluster will get the message (though you must be superuser to override users' protections).

Use wall on a given cluster node when you want the message to go only to users of that particular node.

Cluster-Specific Options of Common Commands

`-c`

This option indicates to certain commands that the command should operate on the whole **c**luster, not just the current machine.

For example,

```
who -c
```

will show you all the current users of all the nodes in the cluster.

The following commands support this option:

- `last(1)` (and `lastb`)
- `users(1)`
- `who(1)`

`-H` *or* `H`

This option means, roughly, "expose **H**idden directories (CDFs) to the operation of the command"—but check the *HP-UX Reference* for the exact meaning in each case.

This option is available with the following commands:

- `chmod(2)`
- `ls(1)` (and `ll`, etc.)
- `tar(1)`
- `test(1)`
- `pwd(1)`
- `fbackup(1M)`
- `find(1)`

`-hidden`

This option is available with the `find(1)` command. It finds all **subfiles** (branches) of context-dependent files, and is particularly useful for backup (see chapter 11 of this manual).

Commands You Cannot Use on a Cluster Client

The following commands will not execute on a cluster client, though they work as normal on the server:

- `fsck(1M)`
- `fsclean(1M)`
- `fsdb(1M)`
- `mkfs(1M)`
- `mount [HFS](1M)`
- `newfs(1M)`
- `tunefs(1M)`
- `umount(2)`
- `update(1M)`

Don't Forget!

The following may trip you up if you don't remember what they do in a cluster:

- System V IPC

  Messages, semaphores and shared memory are *not* distributed in a cluster.
- `/etc/reboot` and `/etc/shutdown`

  These commands take effect on the entire cluster when executed on the server.

Booting a Cluster

To boot a cluster you boot the server, then boot the clients.

Booting the Server

Boot the server just as you would a standalone machine.

For details, refer to chapter 4 of this manual, "Starting and Stopping HP-UX".

Note	For a Series 300 server, make sure you're using the Series 300 version of this manual; for a Series 800, use the Series 800 manual.

Booting a Cluster Client

Before you boot a cluster client, check that all the following conditions are true:

- You have added this machine to the cluster.

  See "Adding Cluster Clients" earlier in this chapter.
- The cluster server has completed its boot sequence and is in multi-user mode.

  You should see a login prompt on the server's console.
- The client is part of only one cluster.

  This means that only one cluster server on the LAN has an entry in its `/etc/clusterconf` file for this cluster node.

  This is not an absolute requirement, but it makes things much simpler.
- No disk attached to the client contains a bootable system.

  This is not an absolute requirement, but it makes things much simpler.

Once you've made sure the client meets all these conditions, you can begin the boot process. This differs depending on whether this is a Series 300 or a Model 815.

Booting a Series 300 Cluster Client

1. Turn on power to the computer.
2. Allow the system to boot in **unattended mode** (that is, do nothing until the boot process is complete).

When the cluster client displays a login prompt, it is ready for use.

Series 300 Client Doubling as Standalone System?

If your Series 300 client will double as a standalone workstation, and therefore has a bootable system on it, you must always boot the client in **attended mode** and choose the operating system you need (as described under "Adding a Series 300 Cluster Client" earlier in this chapter).

Setting up a client this way will complicate the task of administering the cluster: don't do it if you can avoid it.

Booting a Model 815 Cluster Client

First Time Through—Set Primary Path and Autoboot

The first time you boot this machine as a cluster client, you need to set it up to boot automatically over the LAN.

Follow directions labelled "Booting the Cluster Client (Model 815)" under "Adding Cluster Clients" earlier in this chapter.

Booting once Primary Path and Autoboot Are Set.

Always use this method to boot the cluster client once you've initially set primary path and autoboot.

1. Turn power on.
2. Allow the boot to complete uninterrupted.

   Do not press a key when you see the message telling you you have 10 seconds to interrupt.

   (If you press a key by accident, just respond `y` when the system asks if you want to boot from the primary path.)

When you see the login prompt, the cluster client is ready to use.

Model 815 Client Doubling as Standalone System?

If you want your Model 815 client to double as a standalone system, and have a bootable system on a disk drive attached to the 815, you can set set the **alternate boot path** to point to the boot partition on the disk. You do this with the IPL `altpath` command, which works like the `primpath` command shown under "Boot the Cluster Client (Model 815)" in the section "Adding Cluster Clients" earlier in this chapter.

You would still set your primary path and autoboot path to point to the LAN; when you wanted to boot the system on the disk, you would interrupt the autoboot and choose the alternate boot path.

This is a complicated way to set things up: avoid it if you can.

Shutting Down a Cluster

Caution Shutting down the cluster server shuts down the entire cluster.

To shut down the cluster:

1. Log in to the *cluster server* as superuser.
2. Warn users that the system is about to go down, for example:

   ```
   /etc/cwall

   CLUSTER GOING DOWN IMMEDIATELY!!!

   [CTRL] D
   ```

3. Change directories to / (**root**):

   ```
   cd /
   ```

4. Shut down the system:

   ```
   shutdown -h 0
   ```

Shutting Down a Cluster Client

1. Log in to the *cluster client* as superuser.
2. Warn users on this client that the system is about to go down, for example:

   ```
   wall

   SYSTEM GOING DOWN IMMEDIATELY!!!

   [CTRL] D
   ```

3. Change directories to / (`root`):

   ```
   cd /
   ```

4. Shut down the system:

   ```
   shutdown -h 0
   ```

Peripherals: Points to Bear In Mind

The following should help you decide how set up peripherals in your cluster.

Disk Drives	■ File systems must be on cluster server. ■ Cluster server's swap space can be shared by clients. ■ Cluster clients can have their own swap space on local disk.
Tape Drives	■ Supported on both server and clients—but for better performance, do backups on the server unless you have good reason not to. See chapter 11 for examples of backup commands in a cluster.
Printers, Plotters	■ Spooler supported only on the cluster server; shared with clients: print requests can be spooled from any cluster node, but the scheduler runs only on the server. ■ Local printers and plotters can be used as raw devices on clients.
Modems	■ Supported only on cluster server (via `cu(1)` or `uucp(1)`). ■ Not shared.

Reconfiguring the Kernel for a Cluster Node

Each member of a cluster (the server and each cluster client) has its own kernel, which resides in the context-dependent file `/hp-ux`. Chapter 10 of this manual explains when you need to reconfigure an HP-UX kernel, and how to do it.

You may need to modify the cluster server's kernel for any of the reasons described in chapter 10, but if you have to modify a client's kernel, it will probably be for one of these reasons:

- To modify cluster-specific parameters (see Appendix A of this manual and the discussion in the "Cluster Concepts" chapter of *HP-UX System Administration Concepts*).
- To add local swap (see the section "Adding a Local Swap Device" later in this chapter).
- To configure a peripheral local to the client.

Guidelines for Modifying a Cluster Node's Kernel

The procedures for modifying the kernel of a cluster server or cluster client are the same as for a standalone machine, so long as you follow these rules:

- You *MUST* be logged in (directly or remotely) to the cluster node which is to be modified.

  Otherwise you will modify the wrong kernel.

- *Never copy a client's kernel to /SYSBCKUP!*

  Otherwise you will overwrite the backup copy of the *cluster server's* kernel.

- Use SAM wherever possible.

  SAM will manage the kernel, the files associated with it and the backup kernel (`/SYSBCKUP`) so that the server's own files are not overwritten.

  You *cannot* use SAM on an 850 or 855 server. SAM does not support kernel reconfiguration on these machines. Use the procedures prescribed for the 850/55 in chapter 10 of the Series 800 version of this manual, and (if you're adding or reconfiguring peripherals) the manual *HP 9000 Series 800: Configuring HP-UX for Peripherals.*

Determining Swap Requirements

By default (as set up by SAM) each cluster client will swap to the server's swap space, so it's important to make sure that this space will be sufficient. Chapter 10 in the Series 800 version of this manual, and chapter 6 in the Series 300 version, explain how to compute swap space for a standalone system.

To compute the swap space needed by the cluster, you must perform the swap calculations for each machine in the cluster, and then sum the results.

For example, if you have a cluster server and two clients, then your swap space requirement is the sum of the requirements of all three machines. If the server requires 30 megabytes and each of cluster clients requires 20 megabytes, then your total swap space requirement is:

```
20 + 20 + 30 = 70 megabytes.
```

If one of your cluster clients requires a disproportionately large amount of swap space, you should consider attaching a disk to that machine and enabling local swap.

Adding a Local Swap Device

Normally a client will swap to the server's swap space. However, under certain conditions (for example, if you are running large programs from a particular cluster client) it may be more efficient for the client to swap to its own disk, connected directly to it.

You can use SAM to add local swap, or you can do it "manually" by editing files and regenerating the kernel. Unless you are an experienced user and are thoroughly familiar with the process of editing system files and changing the kernel, use the SAM method.

SAM Procedure To Set Up a Local Swap Device

Do this on the CLUSTER CLIENT to which you are adding a local swap device.

1. Make sure that the disk drive which is to be used for local swap does not have any files on it you want to keep.
2. Halt and turn off power to the client.
3. Connect the disk drive physically, following directions in the installation manual for the drive and the peripherals guide for your type of computer (*Configuring HP-UX for Peripherals* for a Series 800; *Installing Peripherals* for a Series 300).
4. Log in to the client as superuser.
5. Get to the `Add a Hard Disk Drive` screen in SAM:

```
SAM Main Menu
      ↓
Peripheral Devices
      ↓
Add a Hard Disk Drive
```

6. Fill in the fields as prompted, using the `Help` key and the peripherals guide for your kind of computer (*Configuring HP-UX for Peripherals* for a Series 800; *Installing Peripherals* for a Series 300) if you need to.
7. Let SAM rebuild the kernel and reboot the system for you.

   The client is now configured for local swap. To see what SAM did to accomplish this, you may want to read through the "Manual Procedure" that follows.

Manual Procedure To Set a Up Local Swap Device (Summary)

1. Edit the `/etc/clusterconf` file to reflect the new swap location.
2. Create a new kernel for local swap:
   a. Modify the kernel source file to reflect the new configuration.
   b. Build and install the new kernel.
3. Connect the disk drive to the cluster client and reboot.

Editing /etc/clusterconf

You need to modify field 5 of `/etc/clusterconf`. You can be logged in on the cluster server or the client to do this.

Suppose your cluster consisted of three nodes (named `cserver`, `client1` and `client2`), then the `/etc/clusterconf` file would look something like this:

```
0800090039dd: #clustercast addresses.  Do not remove.
0800090039dd:1:cserver:r:1:8
080009000565:2:client1:c:1:0
08000900297c:3:client2:c:1:0
```

In each line, the fields are separated by colons. The first five fields have the following meanings:

- Field 1 (up to the first colon) is the **link level address** of the LAN card.
- Field 2 (the number after the first colon) is the **cluster node number.**
- Field 3 is the **cluster node name.**
- Field 4 is the **cluster node type.** It indicates whether this is the cluster server (`r`) or a cluster client (`c`).
- Field 5 (before the last colon) is the **swap server number:** it is the cluster node number of the machine that has the swap device. It will always identify either the cluster server or the current client.

(See the Cluster Concepts chapter of *HP-UX System Administration Concepts* for a complete description of the layout of `/etc/clusterconf`. See also `clusterconf(4)` in the *HP-UX Reference*.)

In this instance, `cserver` is the cluster server (hence the `r` in the fourth field). `cserver` is cluster node number 1, `client1` is cluster node number 2 and `client2` is cluster node number 3.

Both the cluster clients are swapping to the server's disk space: both have a `1` in the fifth field.

To change `client2's` entry so that it swaps to its own local disk, you'd change the `1` in the fifth field of `client2's` entry to a `3`—that is, you make the **swap server number** the same as the **cluster node number:**

```
08000900297c:3:client2:c:3:0
```

Note The remaining steps differ depending whether you are adding a local swap device to a Series 300 client or a Model 815 client. Use the sections labelled "Series 300" or "Model 815" as appropriate

Modifying the Kernel Source File (Series 300)

Caution Do this on the cluster client. Make sure you are *not* on the server.

---

1. Log in as superuser on the cluster client to which you are adding the swap disk.
2. Change directory to `/etc/conf`:

   ```
   cd /etc/conf
   ```

3. Rename the current `dfile`:

   ```
   mv dfile olddfile
   ```

4. Copy the template file `dfile.cnode` to `dfile`

   ```
   cp dfile.cnode dfile
   ```

5. Edit the version of `dfile` you have just created.

   a. Add the driver.

      Enable the appropriate driver by deleting the asterisk (*) in front of `cs80` or `scsi`, depending on which type of disk drive you will swapping to.

      If you are enabling the `cs80` driver, and the disk is connected to a 98625 HP-IB card, delete the asterisk in front of `98625` as well.

   b. Add an entry for swap.

      Find the line

      ```
      * SWAP CONFIGURATION
      ```

      Add the swap entry just before this line. The swap statement is in the form

      `swap` *devname address swplo*

      where:

devname	is the swap device driver (`cs80` or `scsi`);
address	is the minor number of the device (hexadecimal);
swplo	is the swap area location.
	0 = use entire disk for swap
	-1 = use space after file system for swap

      For example,

      ```
      swap cs80 E0000 0
      ```

      indicates swapping to a disk managed by the `cs80` driver at select code 14, bus address 0, using the entire disk space for swap.

Note As a precaution, you may want want to use the -1 option for *swaplo* and configure a very small file system. Then if you should accidentally attach a disk drive that has a file system on it, you will not destroy that file system.

Building and Installing the New Kernel (Series 300)

Do this after you have edited `/etc/conf/dfile` as described above.

1. Make sure you are logged in to the cluster client as superuser.

   The command

   ```
   whoami
   ```

   will tell you. If you are not superuser, log in as superuser now.

2. Change to the `root` directory:

   ```
   cd /
   ```

3. Put the system in single-user mode by typing

   ```
   shutdown 0
   ```

4. Back up the cluster client's kernel by copying it to another name, for example

   ```
   cp /hp-ux /SBCK.cl1
   ```

Caution *Do not* back up the existing kernel by moving it to `/SYSBCKUP`—if you do, you will overwrite the cluster server's backup kernel.

---

5. Change directory to `/etc/conf`:

   ```
   cd /etc/conf
   ```

6. Run the `config` program:

   ```
   /etc/config dfile
   ```

   This creates the files `conf.c` and `config.mk`

7. Build the kernel:

   ```
   make -f config.mk
   ```

8. Move the new kernel to the `/` (`root`) directory:

   ```
   mv hp-ux /hp-ux
   ```

Connecting the Disk Drive and Rebooting (Series 300)

Do this after modifying the kernel source file and rebuilding the kernel as described above.

1. Make sure the disk drive to be added does not have any files on it that you want to keep.
2. Shut down the cluster client:

   ```
   cd /
   shutdown -h 0
   ```

3. Turn off power to the cluster client computer.
4. Install the drive (and HP-IB card 98624 or 98625 if necessary) on the cluster client.
5. Turn on the cluster client computer and reboot.

Prerequisite for the Model 815

The procedure that follows assumes that your 815 is already configured for a disk drive, identified as `disc1 lu 0` in your `S800` file. This is true in the standard configuration (see "Using SAM to Add the Cluster Client (Model 815)", earlier in this chapter).

If your kernel is not configured for a disk drive, you will need to add the appropriate statements to the `io` section (and possibly the driver as well): full directions are in the manual *Configuring HP-UX for Peripherals;* or (for experienced users) there's a quick-reference procedure under "Adding a Peripheral" in chapter 10 of the Series 800 version of the current manual.

Modifying the Kernel Source File (Model 815)

Caution Do this on the cluster client. Make sure you are *not* on the server.

---

1. Log in as superuser on the 815 cluster client.
2. Change directories to `/etc/conf/gen`:

   ```
   cd /etc/conf/gen
   ```

3. Copy the existing `S800` file to a different name, so that you will have a backup copy in case of problems. Give it a name specific to this cluster client. For example:

   ```
   cp S800 S800BCKUP.client2
   ```

4. Edit the S800 file.

   For example:

   ```
   vi S800
   ```

   a. Find the `swap` statement,

      ```
      swap     on     remote
      ```

      It is usually near the beginning of the file.

   b. Change this statement to

      ```
      swap    on     disc1   lu 0  section 2;
      ```

This allocates the entire disk for swap.

Building and Installing the New Kernel (Model 815)

1. Make sure you are on the cluster client whose kernel is to be rebuilt.
2. Regenerate the kernel with `uxgen`, using the edited `S800` file as input:

   ```
   /etc/uxgen S800
   ```

3. Copy the old kernel (`/hp-ux`) in the root (`/`) directory and the old devices file (`/etc/devices`). Give them names specific to this cluster client, and write down the names in case the new kernel does not boot. For example:

   ```
   cp /hp-ux /SBCK.cl1
   cp /etc/devices /etc/DBCK.cl1
   ```

4. Change the working directory:

   ```
   cd /etc/conf/gen/S800
   ```

5. Move `hp-ux` to `/hp-ux` and `devices` to `/etc/devices`, by entering the commands:

   ```
   mv hp-ux /hp-ux
   mv devices /etc/devices
   ```

6. Create the special files (device files) for the new configuration by entering the commands:

   ```
   cd /dev
   /etc/insf
   ```

Connecting the Disk Drive and Rebooting (Model 815)

1. Make sure you are logged in on the cluster client to which you are going to connect the disk drive.
2. Shut down the system:

```
cd /
shutdown -h 0
```

3. Turn off system power.
4. Install the drive.
5. Turn on the system and reboot.

Installing and Removing Programs in a Mixed Cluster

Because Series 800 architecture differs from Series 300 architecture, a program compiled on any Series 300 machine will not run on any Series 800 machine, and *vice versa.*

However, when you write a program to use in the cluster, you will probably want to run it on any machine in the cluster, regardless of architecture, referring to it always by the same name.

This section explains how to install a program so that it can be run from any machine in the cluster, using the same pathname/filename. There is more than one way to do this, but the method recommended here is simple and efficient.

Note If you are installing a third-party package, you will need to find out from the software supplier whether the package's installation procedures will work in the case of a mixed cluster.

Installing a Program File

This is a summary of the recommended procedure; if you've already done this several times, the summary will probably be enough; otherwise, read the summary to get an idea of the tasks, then go on to the section "A Worked Example", which includes explanations and definitions, later in this chapter.

If you have not already done so, you should also read the section "Context Dependent Files" in chapter 12 of the manual *System Administration Concepts*.

Prerequisite

Before you start, you must have an **architecture-specific directory.**

In a mixed cluster, the directories `/bin` and `/usr/bin` are already architecture-specific directories, so if you are planning to install your programs into one of these, you can go straight to step 1.

If you need to create an architecture-specific directory, you can do so using the `makecdf` command.

Note This needs to be done only once, at set-up: you don't have to do it each time you install a program.

For example, the command

```
makecdf -c HP-PA -c HP-MC68010 /usr/local/bin
```

creates an architecture-specific directory `/usr/local/bin` with separate entries for processors of type **MC-68020** and **HP-PA.**

Architecture-specific directories and processor types are explained under "A Worked Example" later in this chapter.

The *HP-UX Reference* contains full details of `makecdf(1M)` and its options and syntax.

Step 1. Compile the program on a Series 800

While logged into one of the Series 800 computers in the cluster, compile the program. For example:

```
cc hello.c -o
```

Step 2. Move the program to its architecture-specific directory

After compiling the Series 800 version of the program, move it to the appropriate architecture-specific directory (see "Prerequisite" above).

To move the file `a.out` from the current directory to the architecture- specific directory `/usr/local/bin` and give the file the permanent name **hello,** issue the command:

```
mv a.out /usr/local/bin/hello
```

Step 3. Repeat on Series 300

Repeat steps 1 and 2 above on one of the Series 300 computers in this cluster.

Removing Architecture-Specific Files and Directories

Note

This is a quick reference—use it only if you already understand what you're doing.

Otherwise, work through the definitions and examples on the following pages, and also read the section "Context Dependent Files" in chapter 12 of *HP-UX System Administration Concepts*.

You'll find a more verbose version of the following examples under "Removing Files—Examples" at the end of the current chapter.

Removing one version

To remove the Series 800-specific or Series 300-specific version of a program file from an architecture-specific directory, go to any machine of the corresponding type in the cluster (a Series 800 cluster node or a Series 300 cluster node) and enter

`rm` *pathname/filename*

For example:

```
rm /usr/local/bin/hello
```

Removing all versions of one type

To remove all the Series 800-specific or Series 300-specific versions of a program file from an architecture-specific directory, go to any machine of the corresponding type in the cluster and enter

`rm` *pathname/**

For example:

```
rm /usr/local/bin/*
```

Removing one branch of an architecture specific-directory

To remove the Series 800-specific or Series 300-specific branch of an architecture-specific directory, go to any machine of the corresponding type in the cluster and enter

`rm -r` *pathname*

For example:

```
rm -r /usr/local/bin
```

Removing an entire architecture-specific directory

If you are quite certain you want to get rid of the *entire* architecture-specific directory (all files of all architectures, and the directory itself), enter

`rm -r` *directory-name*`+`

For example:

```
rm -r /usr/local/bin+
```

A Worked Example

The next few pages provide explanations and a simple example to illustrate the procedure for installing a program onto a mixed cluster, followed by some variations on the theme and examples of *removing* architecture-specific files. If you have performed this procedure several times before, you can probably skip these pages: the summaries on the previous page should suffice.

Some definitions before you start:

These are working definitions to refresh your memory. If you have not already read the "Cluster Concepts" chapter of the *HP-UX System Administration Concepts,* you should stop and do so now, paying particular attention to the section on context-dependent files.

Context

In cluster terminology, a context is an ASCII string which identifies a particular cluster node.

This string appears in each message the cluster node sends over the LAN (every command you type from a cluster node results in such a message).

The context string is built partly from information you supply to SAM when you add the cluster node to the cluster, and partly from information derived automatically at boot time.

Type the command

```
getcontext
```

to see the context string for the cluster node you are currently on.

Processor Type

One of the fields in the **context** string is processor type. There are three processor types:

- **HP-MC68020** for a Series 300 computer that uses MC68020 board.
- **HP-MC68020 HP-MC68030** for any Series 300 model using the MC68030 board.
- **HP-PA** for any Series 800 machine.

CDF

Context-Dependent File. A context-dependent file is a special kind of file which includes entries specific to particular cluster nodes in the cluster. Any file or directory can be made into a CDF (but *DO NOT* modify system files!). That file or directory can then have different contents depending on which cluster node (or group of cluster nodes) is looking at it.

In effect, a context-dependent file (CDF) is a directory that looks like a file, or (in the case of a directory) is a nest of directories that looks like a single directory. For this reason, a CDF is also called a **hidden directory.**

When the server receives a request to use a context-dependent file (or directory), it scans the file for an entry that matches any one of the fields in the context string that forms part of every such request. If it finds a match, it makes the corresponding portion of the file available to the requestor. Thus any given cluster node's view of a CDF may be unique, or a group of cluster nodes (such as all Series 300s) may share the same view. If there is nothing in the requestor's context string that matches any entry in the file, then the file is invisible to that cluster node and you'll get a `not found` message.

Context-Dependent File — See **CDF.**

cdf — See **CDF.**

Architecture-Specific Directory — We are about to create a context-dependent directory, whose contents will differ depending on the architecture of the cluster node that is looking at it. We'll call this an **architecture-specific directory.**

(There are more complete explanations in the "Cluster Concepts" chapter of *HP-UX System Administration Concepts.*)

Task Definition

Let's assume that our cluster consists of an 835 server, an 815 cluster client, and two 319 cluster clients.

We are going to create a directory such that, when one of the Series 800 machines looks at it, it will see only those program files that are executable on a Series 800; and when one of Series 300 machines looks at it, it will see only those program files that are executable on a Series 300. This new directory will be a **context-dependent file,** or **CDF** for short.

Note

Remember that you need do this once at most, not every time you install a new program in the cluster.

If you will always use `/usr/bin`, which SAM converted to a CDF when you configured your mixed cluster, you don't need to create your own architecture-specific directory, though you may want to read through the example to understand what's going on.

After creating the architecture-specific directory, we'll compile and install a single program in two versions, one for the S300 architecture and one for the S800.

Creating an architecture-specific directory

First select the directory where you want compiled code to reside. Let's assume you pick a directory named `/usr/local/bin`.

The directory doesn't have to exist at this point: the `makecdf(1M)` command described below will create it.

The makecdf(1M) command.

In this example, we need to make the directory `/usr/local/bin` into a CDF with separate entries for cluster nodes of processor type **MC-68020** (the 319s) on one hand, and **HP-PA** (the 835 and the 815) on the other.

To do this, we enter the command

```
makecdf -c HP-PA -c HP-MC68020 /usr/local/bin
```

This creates a directory with two branches, only one of which will be visible to any given cluster node. The 815 and 835 will see only program files in the `HP-PA` portion of the file; the 319s will see only those files under `HP-MC68020`.

If I use the `mv` command to move a file into `/usr/local/bin`, it will fall into the `HP-PA` portion of the file if I'm on the 835 or the 815, or into the `HP-MC68020` portion if I'm on one of the 319s.

In your case, the `makecdf(1M)` command string should contain as many processor types as your program needs to take account of, each preceded by `-c` (for context).

Note If you look back at the list of processor types above, you'll notice that **HP-MC68020** appears as part of the context string of *all* Series 300 machines. This means that the directory we've just created will cover most cases—mostly you'll need just one version of a program for the Series 300s in your cluster, and one for the 800s. However, if you will be installing programs designed to take advantage of the floating-point capabilities of a specific Series 300 processor, you will need to include the floating point hardware type in the `makecdf(1M)` command string; and then compile and install those programs on a machine of that type as described in steps 1 and 2 below. See chapter 12 of the *HP-UX System Administration Concepts* for more information.

Now you have a place-holder for each species of compiled code, installing a new program is quite straightforward. Perform the following two steps on one machine of each type of architecture. In this example, we will go through the procedure once on the 835 server and once on one of the 319 cluster clients.

Step 1. Compile the program on Series 800

Let's say we have written a source program named `hello.c` in the directory `/usr/local/csrc`.

In this example, we go to the 835 server and get into the `/usr/local/csrc` directory:

```
cd /usr/local/csrc
```

Now we compile the program:

```
cc hello.c -o
```

This produces the default output file `a.out` in `/usr/local/csrc`.

Step 2. Move the compiled code into the architecture-specific directory.

Still in the `/usr/local/csrc` directory on the Model 835 cluster server, we type

```
mv a.out /usr/local/bin/hello
```

This gives us a program named `hello` in the `/usr/local/bin` directory, just as you would expect if no architecture-specific directory were involved.

The difference is that *only* the Series 800 cluster nodes can see this file. If we now go to one of the 300s and type

```
cd /usr/local/bin

ls
```

there will be no output from the `ls` command: from the point of view of the 319 cluster node, the `/usr/local/bin` directory is empty.

The directory structure under `/usr/local/bin` actually looks like this:

```
            usr/local/bin
                 |
                 |
        -------------------
        |                 |
      HP-PA           HP-MC68020
        |
      hello
```

Architecture-Specific Directory

You can see this structure if you use the `-H` (for **H**idden) option of the `ls` command. In this example, we'd type:

```
ls -H /usr/local/bin
```

The system would respond:

```
/HP-PA /HP-MC68020
```

In effect, `/usr/local/bin` is a directory containing the sub-directories `/HP-PA` and `/HP-MC68020`. The sub-directory `/HP-MC68020` is empty at this point, while `/HP-PA` contains the program file `hello`.

Step 3. Repeat on Series 300.

To complete the installation, we repeat steps 1 and 2 on one of the Series 300 computers.

When you have gone through steps 1 and 2 once on one machine of each type of architecture, the installation is complete. You will be able to type the same program name on any machine in the cluster and the program will run.

In this example, we will *not* need to do anything on the other 319, or on the 815. The 815 and the 835 are code-compatible, so we only need to compile once to cover both of them; and of course the compile we do on one of the 319s will be good for the other.

Why this works

No matter what machine in the cluster I am on when I type

```
/usr/local/bin/hello
```

the right version of the `hello` program will execute. How does the server know which version of the program to run?

Let's say I am on the 815 or the 835. My processor type, which appears in the context string sent to the server every time I type a command, is **HP-PA.** (See above, under "Some definitions before you start", for fuller explanations.)

Software on the server will search the context-dependent directory `/usr/local/bin` for an entry (in effect for a sub-directory name) that matches some part of the context string. It will match on **HP-PA,** and the program in that portion of the directory will run.

Similarly, if I'm on a 319, the server will match **HP-MC68020** in my context string with the corresponding entry in `/usr/local/bin`, and run the Series 300-specific version of the program.

Variations

S800 or S300 only

Suppose you want a given program to run only on the Series 300s in your cluster. Simply follow the steps above, but do them only once, on a 300.

In this example this will create a copy of the program in the `/HP-MC68020` portion of the directory, but no corresponding version in the `/HP-PA` portion.

If you installed the `hello` program this way, then users of any Series 300 machine would be able to run the program, but someone on an 800 attempting to run it would get the message

```
hello: file not found
```

Using a single directory

If for some reason you don't want to make a CDF out the directory where you keep compiled code, you can still create versions of individual program files that are specific to each architecture, though it's more work and is not as tidy a solution.

Example. Let's assume the same cluster as before—an 835 server, and one 815 and two 319s as cluster clients. `/usr/local/bin` is an ordinary directory: you have not made it into a CDF. You have written the program `hello.c` and it resides in `/usr/local/csrc`.

Note The following procedure is not recommended: it's unnecessarily cumbersome. Use the procedure under "Installing a Program File" above unless you have a compelling reason not to.

Step 1. Compile the program

```
cd /usr/local/csrc

cc hello.c -o
```

This gives you a program `a.out` which you can run from the cluster node you are currently on, and from any other cluster node of the same architecture. For the purposes of this example, let's say you are on the 835 server.

Step 2. Create a permanent file

```
mv a.out /usr/local/bin/hello
```

Step 3. Create a CDF

Now you need to set up the file to contain versions of the program specific to the Series 800 on one hand and Series 300 on the other. Use the `makecdf(1M)` command:

```
cd /usr/local/bin

makecdf -c HP-PA -c HP-MC68020 hello
```

This gives you a file `/usr/local/bin/hello` which has entries labelled `HP-PA` and `HP-MC68020`

Step 4. Recompile the program

Now create the Series 300 version. Go to one of the 319s and compile the program as in step 1.

Step 5. Re-install the program

/usr/local/csrc/a.out is now a version of the hello program that will run on a Series 300. You need to keep it in the 300-specific portion of the CDF /usr/local/bin/hello. From the 319 keyboard, type

```
mv a.out /usr/local/bin/hello
```

The program /usr/local/bin/hello will now run from any machine in the cluster.

Removing Files—Examples

Removing one version

Suppose you want to remove the Series 800-specific version of the `hello` program from the context-dependent directory `/usr/local/bin` that we set up in the example under `Creating an architecture-specific directory`.

Go to any Series 800 machine in the cluster and type

```
rm /usr/local/bin/hello
```

This will delete the version of the file that resides in the `/HP-PA` portion of `/usr/local/bin`, but leave the version in the `/HP-MC68020` portion intact.

Removing all versions of one type

Similarly, if you stay on the Series 800 cluster node and type

```
rm /usr/local/bin/*
```

you will remove all the files from the `/HP-PA` portion of `/usr/local/bin` and leave the `/HP-MC68020` portion intact: all the program files that were in `/usr/local/bin` will now be executable only from a Series 300 machine and will be `not found` on any Series 800 machine.

Removing one branch of an architecture specific-directory

Executed from a Series 800 cluster node, the command

```
rm -r /usr/local/bin
```

removes the `HP-PA` portion of the directory, as well as all the files under it, so that `/usr/local/bin` now does not exist from the point of view of any Series 800 machine.

Removing an entire architecture-specific directory

If you were quite certain you wanted to get rid of the *entire* `/usr/local/bin` CDF—all files of all architectures, the structural entries `HP-PA` and `HP-MC68020` and finally the directory itself—then you'd type

```
rm -r /usr/local/bin+
```

In this case it would not matter what machine you were on, but note the plus sign, which acts as an escape character giving you access to the full directory-structure of `/usr/local/bin`.

11

Backing Up and Restoring the System

Users expect you to back up the HP-UX file systems and provide a means to restore those file systems. A backup consists of copying files to a secondary media. A restore consists of copying files from a secondary media to the HP-UX file system. You need to coordinate the two processes (backup/restore). The following table shows types of backup.

Type of Backup	Description
Full, also called Archive	Copying the entire HP-UX file system.
Incremental	Copying files that have changed since the last full or incremental backup. This type usually takes less time than a full backup.
File system	Copying a particular file system (for example, `/etc`, `/users`, or `/usr`).
Directory	Copying files in a directory under a file system according to needs and situations. Typically, you relate this to users, applications, tools, libraries, and utilities.
Other	Backing up particular files as requested.

For each type of backup (and restore), you consider such things as media, commands, and schedules. To help you do effective backups, and subsequently restore them, this chapter discusses planning. Then, it discusses the procedures.

Planning the Backup/Restore Processes

The configuration, size, and complexity of your system should determine your backup/restore procedures. You should not just decide to use a certain procedure. It may not be appropriate.

This module explores the possibilities for backup/restore. Reading the material can help you make decisions about the procedures you want to use. As you read, think about your system and develop a plan for backing up and subsequently restoring your system. If you already understand your alternatives and know what you want to do, skip this module.

Deciding Which Files to Back Up

Back up the entire file system after installing HP-UX or updating to a new release of HP-UX. Then, back up often modified file systems daily. Back up other file systems whenever you change them and cannot afford to lose their files.

Other than performing routine, scheduled backups of certain files, the continual process is to decide what to backup and how often to do it. You always have the options of backing up any of the following entities:

- The HP-UX file system (the root file system).
- A file system under the root (`/users`).
- A directory in some file system (`/users/sue/new-wrk`).
- An individual file (`/util-readme`).

To make these decisions, consider your system and circumstances. For example, the file system on a workstation environment might use 100 Mbytes. You could back up this system on an optical disk in a few minutes. Later, to restore a file, you could mount the file system and recopy the file in minutes. On the other hand, an HP-UX cluster might use more than 400 Mbytes. Backing up the file system to cartridge tapes in an autochanging device takes an hour or more, and restoring a single file can take hours.

The following table provides guidelines for backing up files.

File System	Backup Decision and Schedule
/bin	Possibly never. If you lose this file system, HP-UX is helpless. Your original media has the files. If you make a recovery system, you might back up the files in `/bin`.
/etc	Back up this file system each time you customize its files.
/dev	Back up this file system each time you do something that creates new device files.
/usr	A difficult choice. Back up `/usr/mail` daily. Other backup depends on how often you change files.
/tmp	Possibly never. People should not keep files here that do not exist in another file system.
/users	Daily on most systems. This process becomes more cumbersome as you add users to complex, large systems. You may need an autochanging device. You may need to compress files.
/lost+found	Possibly never. This file system receives files that are orphaned during a file system consistency check. You should return the files to their owners and remove them.
Non-standard file systems	In time, you may create some non-standard file systems (for example, `/net` or a file system for a non-root disk). Back up these file systems as appropriate.
Applications	This depends on your situation. You might backup a robot controller once. An aerospace environment might need continuous backup.
Directories	Users may ask you to backup a certain directory. Accommodate them as appropriate.
Individual files	Whenever you add a file to the system that you cannot afford to lose, back up that file or back up the directory that contains it.

Making Decisions About Media

Options for media include a hard disk, optical disk, magnetic tape, cartridge tape, or flexible disk. Before you do a backup, the corresponding device must have a device file in `/dev` and a device driver in your kernel.

The capacity of the media you select must exceed the size of the file system you back up. You can accommodate limitations by compressing files, backing up smaller parts of a file system, or backing up files to multiple media in an autochanging device.

Optical disks	Accommodate most backups. Mounting a file system on an optical disk lets you use conventional commands to restore files. Capacity equals 320 Mbytes per side. You can use an autochanging device.
Magnetic tapes (9-track tapes)	Accommodate a range of backup situations, especially large systems and UNIX ™ systems. Capacities approximate 100 to 300 Mbytes, depending on the tape length and the bits-per-inch (bpi).
Cartridge tapes	Accommodate typical backup situations, especially the file systems on a standalone workstation. Workstations typically have a cartridge tape drive. Capacities equal 15 Mbytes for a 150-foot cartridge tape, 67.5 Mbytes for a 600-foot cartridge tape. You can use an autochanging device.
Flexible disks	Accommodate small backups (for example, physically transporting files to another user). Capacities range from 360 Kbytes to 1.44 Mbytes. HP does not recommend using flexible disks for typical backups.
Hard disks	Accommodate varied backups, especially archiving documents or code that must often be restored, modified, and backed up again. Match the hard drive to your required capacity.

Selecting Commands (Scripts, Utilities)

Every backup/restore procedure uses at least one command, (script, utility). Since the features of commands (scripts, utilities) vary, you need to study them in detail. Besides using the information presented below, read the entries for the commands (scripts, utilities) in the *HP-UX Reference* manuals.

The fbackup Utility

The `/etc/fbackup` utility backs up files to a raw magnetic 9-track tape. With a pipe to `tcio`, you can back up to a cartridge tape. HP recommends using this utility.

- Use `/etc/fbackup` when you have a secured system that requires you to preserve Access Control List (ACL) information.
- You must use `/etc/frecover` to restore files backed up with `fbackup`.
- You can use `fbackup` for incremental backup, but doing this requires keeping a database of past backups.
- See *fbackup*(1M) and *frecover*(1M) in the *HP-UX Reference* manual to get complete descriptions of these utilities.

The Backup Scripts

A backup script performs a backup while also performing tasks such as checking the file system and creating backup log files. By customizing a script, you can tailor the backup to fit your situation. HP provides the `/etc/backup` script, and you can develop your own scripts.

The `/etc/backup` script performs an archive (full) or incremental backup of a file system to a cartridge tape. The script uses `tcio` and `cpio` (or `ftio`) to do the backup. It also uses other commands to provide additional features. You can use a 9-track magnetic tape by removing `tcio` and redirecting input to `cpio` from the tape device. Study the script and read *backup*(1M) in the *HP-UX Reference* manual to get complete information.

The backup commands continue on the next page.

Commands Used for Backup

/bin/dd — Use *dd*(1) to back up all files on a disk. You cannot back up one of several file systems on a disk. The command works well with an optical disk because the disk usually has sufficient capacity to hold the entire input-file disk. Later, you mount the file system on the optical disk and then recover the files.

```
dd if=/dev/0s0 of=/dev/rscsi bs=1024k
```

Backs up everything on the root hard disk (`/dev/0s0`*) to the disk in the device for* `rscsi`*, using a block size of 1024.*

/usr/bin/ftio — Use *ftio*(1) with 9-track magnetic tape for faster I/O than you get using *cpio*(1). You can adapt the command for use with cartridge tapes. The `backup` script lets you use `cpio` or `ftio`, depending on how you customize it. See the *HP-UX Reference* manual if you use *ftio*(1) because it has many features.

/usr/bin/tar — This long-standing command is found on most UNIX ™ systems. You can use *tar*(1) to archive file systems, directories, or individual files on a cartridge tape or flexible disk. You typically do this to physically transfer files to another person. You can transfer files archived by `tar` among UNIX ™ systems. The command handles tape errors ungracefully; so it is not recommended for regular backups.

```
tar -cvf /dev/rct*
```

Creates an archive of all files in the current directory on the tape in the named device.

```
tar -xvf /dev/rct
```

Recovers all files on the tape in the named device in the current directory.

Commands Used With Pipes for Backup

/bin/find — Use *find*(1) as the command that seeks files. If you have an HP-UX cluster, using the `-hidden -depth` options preserves the Context Dependent Files (CDFs). This is important because some commands to not back up the hidden directories.

/bin/cpio — Use *cpio*(1) as the command that does the copying. Unless you copy everything (the default), `cpio` requires file specification, and this specification can be a pattern that requires single quotes. You must specify files exactly in terms of punctuation and pathnames. See the *sh*(1) entry in the *HP-UX Reference* manual if you need help with specifying patterns.

/usr/bin/tcio — Use *tcio*(1) as the command that operates on the character device file for a cartridge tape (for example, `/dev/update.src`).

You can use `find` and `cpio` without `tcio` and with redirection to a file when using media other than a cartridge tape.

The following example, which you can type as one line, shows using the commands with pipes to send all files in `/tmp` to a cartridge tape:

`find . -hidden -depth -print |`	*Finds all files including CDFs in the current directory.*
`cpio -ovc '/tmp/*' |`	*The '/tmp/*' is the pattern for copying. You need single quotes because the pattern has a wildcard character.*
`tcio -o /dev/update.src`	*Outputs to the media in a cartridge tape drive*

The *HP-UX Reference* manuals have additional information about these commands. The procedures explained later assume you will read the entries and get detailed information before you do a backup or restore.

Matching the Backup Media to System Size

The capacity of the device selected for backup must exceed the size of the file system you back up. Capacity increases as a system increases in size, number of users, and number of supported systems. To choose a backup media and its device, examine your file systems and available media.

Use `du`, `df`, and `bdf` to obtain information about your system size and compare the information they display with the known capacities of available media. For example, if `bdf` shows the `/users` directory occupies *46%* of a disk having a capacity of 590 Mbytes, you could use a 9-track magnetic tape or an optical disk to back up that directory. You might be able to use a 600-foot cartridge tape (capacity = 67.5 Mbytes) if compressing the files makes them small enough. Making these decisions illustrates the planning process. If you have insufficient capacity, use the following techniques:

- Back up smaller pieces of the system (for example, a single file system such as `/usr`). In a file system such as `/users`, you could back up an individual user.
- Compress the files you back up; uncompress the files you restore.
- Use an autochanging device that has cartridge tapes or optical disks.

Deciding on Manual or Automatic Backups

You can perform a backup manually or use `cron` (the clock daemon) to perform a backup automatically.

If you decide to back up files automatically, you may have some limitations (for example, having a secured system can limit your use of commands). In most cases, you should perform periodic backups automatically and perform random backups manually (directly). For example, you might back up `/users` automatically on a daily basis and back up `/etc` after updates or customizations.

If you decide to automate parts of your total backup, use the `/etc/fbackup` utility or use a backup script (for example, `/etc/backup`). You can develop your own script and use any appropriate combination of commands if you have shell programming expertise.

This manual does not provide much information on automatic backups because situations can vary widely. Also, the manual does not explain any shell programming. Instead, the manual assumes you can study the commands and adapt their use in scripts to your situation.

Moving on to an Actual Backup

Once you plan your backup, have a grasp of your situation, and understand the commands you plan to use, move on to the next module to see the procedures. Use the following guidelines:

- Set up your backup and restore procedure after you install or update the system, make a recovery system, and perform the initial customizations.
- The planning process may lead you to install additional mass-storage devices (for example, an optical disk drive). You may also need to obtain additional media (for example, a 320-Mbyte optical disk).
- Plan for backup and restore in terms of your immediate and projected needs. For the long term, use maximized estimates of your needs.
- Try to get an effective match among type of system, use of media, and procedure (the scripts or commands you run directly or automatically).

Procedures for Manual Backup and Restore

This module explains manual backup and restore procedures. Any procedure can be done manually, and many can be automated (run by `cron`). You see the manual procedures first; later, you see an automated backup process.

Prerequisites and Conditions

- Do the planning described in the previous module. Backing up files is not a trivial task, and restoring files depends on the backup method.
- If you have a secured system, you must use `fbackup` for backup and use `frecover` for recovery because these commands preserve Access Control List (ACL) information.
- If you use a backup script, the script resides in `/etc` after an installation, or in `/etc/newconfig` after an update. After an update, copy the scripts you use into `/etc`.
- Any backup log files usually are put in `/etc`.
- Match the commands to your media (for example, use `tcio` when you use a cartridge tape).
- For commands such as `cpio`, specify **patterns** according to conventions defined in the *sh*(1) entry in the *HP-UX Reference* manual.

`cpio [options] *`	*Specifies all files in current directory, the default.*
`cpio [options] users/bob/stuff`	*Restores stuff from /users/bob.*
`cpio [options] 'usr/spool/lp/*'`	*Requires single quotes because the * makes the specification a pattern.*

Specify patterns exactly because the punctuation and pathnames are critical.

The procedures you can use begin on the next page.

Manual Backup Using the Fbackup and Frecover Facilities

The /etc/fbackup and /etc/frecover facilities provide comprehensive ways to back up a system and subsequently recover files from a 9-track tape (reel-to-reel tape).

Prerequisites and Conditions

- You must use these facilities for backup and recover if you have a secured system and need to preserve ACL information.
- Read the *fbackup*(1M) and *frecover*(1M) entries in the *HP-UX Reference* manual before you use the facilities.
- Install required devices and make the corresponding device files.
- Compare the capacities of your backup media with the size of your system. Make any necessary adjustments in your procedure (for example, backing up the system versus backing up a file system).

Backup Procedure

HP cannot recommend an exact procedure. Study the facilities, noting their features. Then, use them according to your needs. The examples show backing up directories, subdirectories, and files as implied recursively in the text file named /usr/adm/fbackupfiles/graphfile to a 9-track tape at /dev/rmt/0h. Type the entry as one line.

- A full backup:

```
/etc/fbackup -u0f /dev/rmt/0h -g /usr/adm/fbackupfiles/graphfile -c
/usr/adm/fbackupfiles/fb_config
```

- A weekly backup:

```
/etc/fbackup -u3f /dev/rmt/0h -g /usr/adm/fbackupfiles/graphfile -c
/usr/adm/fbackupfiles/fb_config
```

- A daily backup:

```
/etc/fbackup -u8f /dev/rmt/0h -g /usr/adm/fbackupfiles/graphfile -c
/usr/adm/fbackupfiles/fb_config
```

Restore Procedure

Having backed up files using the examples shown for the previous procedure, you can recover the files. As stated earlier, see the *frecover*(1M) entry in the *HP-UX Reference* manual to get information.

The following command recovers all files on the backup media and organizes them into the directories from which they were backed up:

```
/etc/frecover -r -f/dev/rmt/0h
```

The following command recovers all files specified in `/usr/adm/fbackupfiles/graphfile` and executes a shell script listed in `/usr/adm/fbackupfiles/fr_config` if the recovery process has an error:

```
/etc/frecover -x -g /usr/adm/fbackupfiles/graphfile -c
/usr/adm/fbackupfiles/fr_config -f/dev/rmt/0h
```

Manual Backup Using dd

The `dd` command (copy input file to output file with conversions) backs up all files on a disk.

Prerequisites and Conditions

- See the *dd*(1) entry in the *HP-UX Reference* to get complete information.
- The command lets you convert ASCII files to EBCDIC files, and conversely.
- Do not use the command to back up file systems having CDFs or ACL information.
- Have adequate capacity on an optical disk or 9-track magnetic tape. Do not use the command with a cartridge tape.

Backup Procedure

The following command copies all files for the input file (the typical device file for the root directory) to the output file (the typical device file for the optical device) with a block size of 1024 Kbytes.

```
date; time; dd if=/dev/rdsk/0s0 of=/dev/rscsi bs=1024k
```

Along with date and time, archives all files in the root directory

Restore Procedure

1. Insert the disk into the optical device.
2. Make a directory for receiving the files (for example, `/trash`).
3. Mount the directory (for example, `mount /trash /dev/rscsi`).
4. Examine the directory, using typical commands. Find the files you want to restore, and copy them to your HP-UX system.

Manual Backup Using find, cpio, and tcio

You can adapt the `find`, `cpio`, and `tcio` commands to many backup situations, especially if you back up file systems having CDFs.

Prerequisites and Conditions

- The *HP-UX Reference* manuals have information about the commands, options, arguments, and patterns.
- Use a backup cartridge tape that has sufficient capacity:
  - A 600-foot CT = 67.5 Mbytes.
  - A 150-foot CT = 16 Mbytes.
- When your need exceeds the capacity, backup as follows:
  - Compress the files (the examples show this).
  - Use an autochanging device for cartridge tapes.
  - Back up a selected file system to a cartridge tape.
- Have a device file for each backup device.
- You can use `find` and `cpio` to back up files to other media. Do not use `tcio`; instead, redirect the output from `cpio` to the device file for the media. The later section named "Manual Backup to a Cartridge or 9-track Tape" shows the differences in the command lines. The later section named "Manual Backup to an Optical Disk" shows how to use `find` and `cpio` with optical media.

Backup Procedure

1. Change to a directory, for example `/users`:

   cd *directory*

2. Execute the following command (the device file might be `/dev/update.src`):

   find . -hidden -depth -print | cpio -ocv | compress | tcio -o /dev/*dev_file*

Restore Procedure

1. Insert the media in the appropriate device. Wait until it is ready for operation.
2. Assume you want to restore files as follows:
   a. Get all files for `denise` in `/users`.
   b. Uncompress the compressed files.
   c. Put the restored files in `/tmp`.
3. Execute the following command:

   ```
   cd /tmp
   tcio -i /dev/update.src | uncompress | cpio -icdv 'users/denise/*'
   ```

   *You need quotes because the specification has a * character.*

   To recover the files in the background and log off, execute:

   ```
   cd /tmp
   nohup "tcio -i /dev/update.src | uncompress |
   cpio -icdv 'users/denise/*'" &
   ```

4. Wait for the process to complete, which can take considerable time.

Manual Backup Using tar

The `tar` command (tape archive) provides a traditional means to back up and restore files. For example, most UNIX ™ systems can run `tar`.

Prerequisites and Conditions

- You have an uncomplicated backup situation.
- You have a cartridge tape drive and character device file (the procedure assumes `/dev/update.src`).
- The backup can fit on a 150- or 600-foot cartridge tape.
- `tar` handles tape errors ungracefully, so do not use the command for critical backups.
- Not a recommended procedure, especially for Series 800 systems.

Backup Procedure

`cd` *directory*	*Change to a directory (*`/users/work`*).*
`tar -cvf /dev/update.src*`	*Archives files in the directory to the tape.*

Restore Procedure

1. If necessary, make a directory to receive the backed up files. For example:

   ```
   mkdir /tmp/trash
   ```

2. Execute the following commands:

`cd` *directory*	*Change to a directory (*`/tmp/trash`*).*
`tar -xvf /dev/update.src`	*Restores files on tape to the directory.*

Manual Backup Using the Backup Script

The `/etc/backup` script lets you back up file systems or all files on a disk. You can do a full or incremental backup. The script performs tasks related to backup. You can customize the backup by editing the script.

Prerequisites and Conditions

- Match the backup of file systems to the capacities of your media: 16 Mbytes for a 150-foot cartridge tape, 67.5 Mbytes for a 600-foot cartridge tape, approximately 120 to 300 Mbytes for a 9-track magnetic tape, depending on length and bits-per-inch density.
- With some editing, you can use `cpio` or `ftio` to do the copying.
- The procedure assumes ability to edit shell scripts.

Backup Procedure

1. Become the root user. If necessary, change to the root directory.
2. The file systems you want to back up should be mounted, but you want the system to be quiet. The root file system (`/dev/dsk/0s0`) is always mounted. In most cases, run the backup script when you know people will not be using the system.
3. If necessary, insert the backup media in the backup device.

   The procedure continues on the next page.

4. Do an archival or incremental backup as follows:
   a. For an archival backup, execute: `/etc/backup -archive`
   b. For an incremental backup, execute: `/etc/backup`
5. The system verifies the backup device, possibly `/dev/update.src` for a cartridge tape device. Specify a character device file. You can get error messages during this step:
   a. `Warning: you may be backing up to a file.` Exit the backup procedure using Ctrl-D. Restart the process, and next time, specify a character device file.
   b. `errno: 20, Can't write output` (plus another line). You ran out of media. Remove the tape, insert another tape, and press Return to continue. The `tcio` command produces a different wording, but you to the same thing.
   c. Ignore the above messages if you have an autochanging device.
   d. `Optimal acl entries for ...` . Some files have ACL information. Use the `fbackup` command if you get this message.
6. When the backup finishes, remove the media, label it, and store it in a safe place. For an autochanging device, note the order of the tapes.
7. Unmount any mounted file systems using `umount -a`.
8. Use `fsck -p` to check all file systems, or for a first file system mounted on `/dev/dsk/1s0` (for example), use `fsck -p /dev/rdsk/1s0`. Check all systems touched by the backup.
9. Return to the normal run level by using `reboot`.
10. To see backup messages, examine `/etc/backuplog`.

The following pages have the restore procedures.

Restore Procedure When Using the Backup Script

If you use `/etc/backup` to back up files on a flexible disk, cartridge tape, or 9-track magnetic tape, you can use the following procedures to restore files. Use a table of contents listing procedure to determine how to specify the pathnames (slash or period-slash).

If the backup pathname was ...	The restore specification is ...
./path	path
/path	/path

To List the TOC for a 9-track Tape or Flexible Disk.

1. Execute:

   `cpio -ictv < /dev/`*raw_device_file_name*

2. You see a displayed list. Note the initial part of the pathnames. Use the appropriate specification during the restore process.

To List the TOC for a Cartridge Tape.

1. Execute:

   `tcio -i /dev/`*raw_device_file_name* `| cpio -ictv`

2. You see a displayed list. Note the initial part of the pathnames. Use the appropriate specification during the restore process.

Restore procedures for 9-track tapes and flexible disks appear on the next page.

Restore Procedure for 9-track Tapes or Flexible Disks

1. You need not be in the single-user state.
2. Insert the media into the appropriate device and wait until it is ready.
3. Change to the root directory (`cd /`).
4. To restore files, proceed according to your case:
   a. A single file such as `./users/joe/report` for a device file such as `/dev/tape`, execute:

      ```
      cpio -icBdmuvx users/joe/report < /dev/tape
      ```

   b. A directory such as `./users` for a device file such as `/dev/tape`, execute:

      ```
      cpio -icBdmuvx users/\* < /dev/tape
      ```

   c. A full archive for a device file such as `/dev/tape`, execute:

      ```
      cpio -icBdmuvx < /dev/tape
      ```

Restore Procedure for Cartridge Tapes

1. You need not be in the single-user state.
2. Insert the cartridge tape into the appropriate device and wait until it is ready.
3. Change to the root directory (`cd /`).
4. To restore files, proceed according to your case:
   a. A single file such as `./users/joe/report` for a device file such as `/dev/update.src`, execute:

      ```
      tcio -i /dev/update.src | cpio -icdmuvx users/joe/report
      ```

   b. A directory such as `./users` for a device file such as `/dev/rct`, execute:

      ```
      tcio -i /dev/rct | cpio -icdmuvx users/\*
      ```

   c. A full archive for a device file such as `/dev/update.src`, execute:

      ```
      tcio -i /dev/update.src | cpio -icdmuvx
      ```

Manual Backup to a Cartridge or 9-track Tape

Earlier, you read about backup procedures that used certain commands. This module repeats this information to a degree, but it focuses on tapes and has some different examples.

Assuming you installed your system from cartridge tape, you have a tape drive and that drive has a character special device file, probably called `/dev/update.src`. Having these things makes it convenient to back up selected files or a whole file system to a cartridge tape.

The same logic applies if you installed your system from a magnetic tape (also known as 9-track tape) that is used in reel-to-reel tape drives. The tape drive also has a character special device file, probably called `/dev/tape`.

Prerequisites and Conditions

1. In the examples below, `cpio` overwrites everything currently on your tape with the new information. This has two implications:
   a. Do *NOT* use tapes having information you want to save.
   b. Change to the directory having the files you want to back up before you begin the backup process.

      `cd` *directory_name*

2. You must know the device file name. Besides `/dev/update.src` (for cartridge tapes) and `/dev/tape` (for 9-track tapes), system administrators may use other names for device files for cartridge and tape drives.

The following pages describe the procedures, cartridge tapes first, and magnetic tapes second.

Backup to a Cartridge Tape

1. Insert the cartridge tape and wait for it to load.
2. Have an appropriate character device file (for example, `/dev/update.scr`).
3. Change to the correct directory (for example, `/users/project/new-work`).
4. To back up all files and subdirectories, execute:

   ```
   find . -hidden -depth -print | cpio -ocxv | tcio -o  /dev/update.src
   ```

   To back up only the files in the directory, execute:

   ```
   ls | cpio -ocx | tcio -o /dev/update.src
   ```

   To back up selected files, execute:

   `ls` *file_name file_name* ... `| cpio -ocBx | tcio -c /dev/update.src`

   Separate file names by blank spaces. You can specify a relative path and name for *file_name*.

Restore from a Cartridge Tape

1. Change to the directory receiving the files (for example, `/tmp`).
2. Insert the cartridge tape and wait for it to load.
3. Recover files using the examples to suggest specification of files:

`tcio -i /dev/update.src | cpio -icdmux`	*Recovers all files on the tape, including special files.*
`tcio -i /dev/update.src | cpio -icdm users/sally/junk`	*Recovers junk in /users/sally*
`tcio -i /dev/update.src | cpio -icdmux 'dev/*'`	*Recovers all device files and subdirectories.*

Backup to a Magnetic Tape

Besides using a backup script or the `fbackup` utility, which were described earlier, the following procedure shows one way to back up to a magnetic tape:

1. Place the reel on the drive and load it.
2. Have an appropriate character device file (for example, `/dev/tape`).
3. Be in the correct directory.
4. To back up all files and subdirectories, execute:

   ```
   find . -hidden -depth -print | cpio -ocBx > /dev/tape
   ```

   To back up only the files, execute:

   ```
   ls | cpio -ocBx > /dev/tape
   ```

   To back up selected files, execute:

   `ls` *file_name file_name* ... `| cpio -ocBx>/dev/tape`

   Separate file names by blank spaces. You can specify a relative path and name for *file_name*.

Restore from a Magnetic Tape

1. Change to the directory you want to receive the files.
2. Insert the magnetic tape and wait for it to load.
3. Execute:

   `cpio -iBcdmux [`*patterns*`] /dev/tape`

   In the example, *patterns* specifies the files to receive enclosed in single quotes (for example, `'/users/sue/*'` to restore all the files in Sue's directory under `/users`. Not specifying a pattern retrieves all files on the tape.

Manual Backup to an Optical Disk

An optical disk provides an ideal media for backing up a file system or an entire disk.

- Each disk has two sides, and each side has a capacity of 320 Mbytes.
- Thirty-two optical disks can be used in an autochanging device. This provides a backup capacity to approximately 20 Gbytes.
- You can mount a file system on an optical disk and use conventional commands to examine and restore files.

Prerequisites and Conditions

1. Your system must have a S6300 Model 650/A optical disk drive or a Series 6300 Model 20GB/A optical disk library system. Both devices connect to a SCSI adapter card.
2. To install and set up the Model 650/A, use the documentation for the optical disk drive, the adapter card, and the documentation for installing peripherals for your series.
3. To install and set up the Model 20GB/A, HP recommends that you contact your support center. Have them install the device and make required device files.
4. For the Model 650/A, ensure that `/etc/conf/dfile` has a SCSI driver named `scsi`. For the Model 20GB/A, the configuration description file needs the SCSI driver named `ac`. If the file does not have these drivers, reconfigure the kernel using the procedures in the chapter named "Reconfiguring the Kernel".
5. Make block and character (raw) device files for the optical disk. These device files work for Sides A and B. The examples assume Major Number 7 for a block device and 47 for a character device, Select Code 17 (set on the SCSI adapter card), Bus Address 0 (set on the optical drive), and Unit Number 7 (the default SCSI address on the SCSI adapter card):

```
mknod /dev/bscsi b 7 0x110007
mknod /dev/rscsi c 47 0x110007
```

6. Set the ownership and permissions. Device files associated with mountable file systems must have read/write restrictions for group and other.

```
chown root /dev/bscsi /dev/rscsi
chmod 600 /dev/bscsi
chmod 600 /dev/rscsi
```

7. Initialize Side A of the optical disk using the raw device file:

```
mediainit /dev/rscsi
```

   To use Side B, you must also initialize that side.

8. You will have a file system to back up (for example, `/users`). Or you may back up the entire root file system (/). Beyond this, you need a directory for the backup file system you intend to mount on the optical disk. The procedure shown below assumes a backup target directory named `target` that you will mount on the optical disk. Use `mkdir` to create `/target`.

9. To make it possible to mount a file system on the optical disk, you first create a file system on Side A of the optical disk. The `/etc/disktab` file must contain the entry for the optical disk used in the example (or it must have an equivalent entry).

```
newfs /dev/rscsi hpS6300.350A_noswap
```

   The example shows the command, device file, and entry from disktab

   To mount a file system on Side B, you must also create a file system on that side.

10. You have completed the tasks required to use one side of an optical disk for backup (both sides if you initialized Side B and created a file system on it).

Backup to One Side of an Optical Disk

The procedure assumes a file system you want to back up named `/srce` (for backup source). The procedure assumes you back up the source file system to a file system named `/target`, which is mounted on the optical disk. If you want to automate the backup process, you could create a script that performs the tasks shown below and work through the section that describes automating a backup.

1. As the root user, insert an optical disk with Side A up.
2. If necessary, change to the source directory:

   ```
   cd /srce
   ```

3. Mount the target directory on the optical disk so it can receive the backup. Use the block device file.

   ```
   mount /dev/bscsi /target
   ```

4. Execute the following command to perform the backup. Including the `-hidden` and `-depth` options backs up the CDFs if you have an HP-UX cluster. If you do not have a cluster, you can omit them.

   ```
   find . -hidden -depth -print | cpio -pdmuvx /target
   ```

   Wait for the backup to complete (a few seconds to several minutes, depending on the size of the file system).

5. Change to the root directory and unmount the optical disk. Do not remove an optical disk without unmounting it:

   ```
   cd /
   umount /dev/bscsi
   ```

6. Store the disk in a safe place. If you automate the backup, you manually insert and subsequently remove the disk according to your backup schedule.

Restore from an Optical Disk

1. Assuming the situation described above, insert Side A of the optical disk that has the files to be restored.
2. Mount a file system on the optical disk using the block device file:

   ```
   mount /dev/bscsi /target
   ```

3. Change to the target file system (the system that received the backup files):

   ```
   cd /target
   ```

   You may want to list the files to verify that they exist in this directory.
4. At this point, you have some options:
   a. You can restore a particular file to the original directory.

      ```
      cp /target/stuff /srce
      ```

   b. You can restore files that fit a specified pattern to the temporary directory.

      ```
      cp /target/*.rpt /tmp
      ```

      Then, you could restore these files to their owners as requested.
   c. You can restore all the files on the backup media to a specified directory (typically `/tmp`). Or you could restore these files to the original source directory (this can cause problems related to overwriting files). The example assumes you have a reason to restore the backed up files to the original directory.

      ```
      find . -hidden -depth -print | cpio -pdmuvx /srce
      ```

      In short, use conventional methods to restore a file or files.
5. When you finish restoring files, change to the root directory, unmount the file system on the optical disk, remove the disk, and store it in a safe place:

   ```
   cd /
   umount /dev/bscsi
   ```

Doing Automatic Backups

The backup procedures mentioned so far assume you use a command (or script) directly. You can automate the backup process via *cron*(1M) by adding a line to `/etc/lib/cron/crontabs/root`. Except for `fbackup`, which is a compiled command, you can edit the backup script run by `cron` so it backs up the system according to your needs.

Prerequisites and Conditions

- You can use `fbackup`, `backup`, or your own script to do do an automated backup. Automating the process does not restrict your use of commands, scripts, or utilities.
- Use `fbackup` when you have a secured system.
- Using the `backup` script with the `-hidden`, `-depth`, and `-print` options in the `find` command works well when you have an HP-UX cluster.
- If you do not know the relationships among *cron*(1M), *crontab*(1), and *at*(1), examine their entries in the *HP-UX Reference* manual.
- To perform an automatic backup, you edit a file named `root`, which resides in the `/usr/spool/cron/crontabs` directory.
- When doing an automatic back up with `cron`, choose a time when no one will be working. Send users a message requesting that they log off.
- The system clock must be accurate if you use `cron`. The current time and date should be accurate as should the time and date in any file such as `/etc/archivedate`.

- Use tapes as follows:
  - One cartridge tape, having sufficient capacity, inserted in a tape drive such as an HP 9144. A 150-foot tape backs up 16 Mbytes, a 600-foot tape backs up 67.5 Mbytes.
  - Multiple cartridge tapes inserted in an autochanging tape drive such as an HP 35401. Each tape must have sufficient capacity to back up the specified file system. If you use `backup` with, for example, four cartridge tapes in an autochanging device, the script must have the following items:

    ```
    cpio -ocx  |
    tcio -o -l 1 -n 4 $dest
    ```

  - One reel of 9-track magnetic tape in a drive such as an HP 7980A. At 6250 bits-per-inch, the tape holds about 140 Mbytes. If you use `backup` with this device, you must change the following items:

    ```
    cpio -ocx
    tcio -o $dest
    ```

    to this item:

    ```
    cpio -ocBx > $dest
    ```

- Use disks as follows:
  - One optical disk inserted in an optical disk drive such as an HP 650A. A 2-sided optical disk backs up 320 Mbytes per side.
  - An HP hard disk drive having the required capacity.
  - Do *NOT* use flexible disks to back up HP-UX. You could use flexible disks to back up a particular file system (for example, your current project).

The procedure for automating via `fbackup` *begins on the next page.*

Using fbackup for Automatic Backup

The command accommodates a backup to a 9-track magnetic tape in varied situations and capacities. It also accommodates a backup of secured systems because it preserves ACL information.

Prerequisites and Conditions

- Have all necessary devices and device files.
- Have adequate media for your file systems.
- Study the *fbackup*(1M) entry in the *HP-UX Reference* manual before you begin. The command has many features and options.
- The entry in the crontabs file named `root` has six fields: *minute, hour, date, month, day_of_week,* and *command executed by shell.* An asterisk specifies all values for a field.

Backup Procedure

1. Locate your existing cron file. If you do not know the existing cron-file name, create a copy of it by executing:

   `crontab -l >` *new_cron_file* *You might use a name such as* `croninfo`.

2. Add a line to your existing cron information file to automatically start the back up procedure. For example, to automatically perform incremental (verses the full archival) backup at 11:55 p.m., add the following line:

   ```
   55 23 * * 1-5 /etc/fbackup
   ```

   The procedure continues on the next page.

3. Use the `crontab` command with the file's name as an argument. For example, if you named the cron information file `croninfo`, enter:

   ```
   crontab croninfo
   ```

   This creates a file named `/usr/spool/cron/crontabs/root` (which overwrites the existing `/usr/spool/cron/crontabs/root`). Examine this file to see that it will perform the backup as scheduled.

4. If you run a backup from `cron` as described so far, do the following things each morning:

   a. Examine the information and messages listed to the printer during the previous night's backup process.

   b. Remove, label, and store the backup medium.

   c. Install the next backup medium in your series of media.

Restore Procedure

To restore files backed up using `fbackup`, use the `frecover` command, which has options similar to `fbackup`. If you want, you can define a `graphfile` that indicates which files should be recovered and a `config` file to customize the behavior of `frecover`.

For example, the following command, which must be typed as one line, recovers all files on media for a specified device file according to specifications placed in `/usr/adm/fbackupfiles/graphfile` and executes the shell script specified by `/usr/adm/fbackupfiles/config` if the recovery process encounters an error.

```
tcio -i /dev/special_dev_file | /etc/frecover -x -f - -g
/usr/adm/fbackupfiles/graphfile -c /usr/adm/fbackupfiles/config
```

Restoring Files after a System Crash

All the procedures for backup and restore mentioned so far assume HP-UX is working. If you have a system crash, the restore procedure changes dramatically. The procedure for restoring your system depends on many factors. For example:

- How bad was the crash? Did you lose everything? Can you start HP-UX? Did you lose date files and system files? You can encounter many variations of this scenario.
- Do you have a **recovery system**? The chapter on "Constructing an HP-UX System" explains how to make a recovery system. You had to make this system, and if you did not make it, you have limited your options for restoring the system.
- Do you still have your original media, the tape containing the HP-UX filesets? If everything else fails, you can always reinstall HP-UX and restore whatever files you can obtain for commands, users, and applications from backup media.

Prerequisites and Conditions

- You do not have any real prerequisites. You may have noticed some strange behavior before the system crashed. If you did, try to recall what was happening.
- As for the condition of the system, you might have anything from a system that will not start up to a system you cannot stop expect by turning off the computer. You may have a hung port or a malfunctioning getty.

Manual Procedure for Recovering from a System Crash

1. If you have figured out what caused the crash and you can startup HP-UX, work through your directories, restoring missing files by any appropriate means. While this statement is simple enough, the process may take a long time, and you may need to explore the system at considerable lengths.
2. If the system will not start up and you have a recovery system, use the procedures described in the chapter named "Managing the File System" under recovery systems to get the system going again.

3. If the system will not start up and you have no recovery system, you must reinstall HP-UX from your original media and get it going again for users, applications, and so on from backup media.

A

Operating System Parameters

If you need to customize your HP-UX kernel, you can use these tunable parameters with the `config` command or the SAM tool. These programs are described in *System Administration Tasks*.

If you change these parameters you can adversely affect the operation of HP-UX. *Be sure you know all the implications of using the parameters before you configure your system.*

Operating system parameter values may change from release to release because new functionality often requires tuning the kernel. Please be aware of this.

Logical Groups of Paremeters

Although this reference appendix is organized alphabetically, the following table shows several logical groups of operating system parameters.

Group	List of Parameters
Message	msgmap, msgmax, msgmnb, msgmni, msgseg, msgssz, msgtql, mesg
Semaphore	semaem, semmap, semmni, semmns, semmnu, semume, semvmx, sema
Shared Memory	shmseg, shmmni, shmmax, shmbrk , shmmaxadder, shmmin, shmall, shmem
Accounting Code	acctresume, acctsuspend, timeslice
File System	nfile, nflocks, ninode, ntext, nbuf
Process	nproc, maxuprc, maxdsiz, maxssiz, maxtsiz
Networking	netmeminit, netmemmax, netmemthresh, netisr_priority
Cluster	dskless_fsbufs, dskless_cbufs, dskless_mbufs, dskless_node, maxswapchunks, minswapchunks, ngcsp, num_cnodes, selftest_period, server_node, serving_array_size, using_array_size
Hardware	fpa, num_lan_cards, parity_option, num_pdn0
MS-DOS	dos_mem_byte
Miscellaneous	maxusers, scroll_lines, dst, timezone, unlockable_mem, npty, ncallout, argdevnblk, ndilbuffers

acctresume

Name

`acctresume`—resume accounting due to disk usage

Range

-100 to 101

Default

4

Use

The system automatically disables process accounting when the available space on the file system where the accounting file resides falls below a certain threshold. The threshold is described under `acctsuspend`. The system also automatically re-enables process accounting when sufficient space becomes available. The parameter, `acctresume`, is the threshold (percentage of free space) which the system must have to re-enable process accounting. This percentage is added to minimum free percentage (minfree) for the file system.

A value of zero re-enables accounting when the free space reaches minfree. A value less than zero allows process accounting to use the space which is reserved for superuser use. A value greater than 100 prevents process accounting from ever being re-enabled because of available space.

When accounting is re-enabled in this way, the message:

```
"Accounting resumed"
```

is printed on the system console.

`acctresume` is only relevant to systems which turn on process accounting.

Cost

None.

Dependencies (interactions with other system values)

`acctsuspend < acctresume`

acctsuspend

Name

acctsuspend—suspend accounting due to disk usage

Range

-100 to 100

Default

2

Use

The system automatically disables process accounting when the available space on the file system where the accounting file resides falls below a certain threshold. The parameter, acctsuspend (specified as a percentage of free space), is the threshold. This percentage is added to minimum free percentage (minfree) for the file system.

A value of zero disables accounting when the free space falls below minfree. A value less than zero allows process accounting to use the space which is reserved for superuser use. If the sum of acctsuspend and minfree is less than zero, process accounting can never be disabled for this reason.

When accounting is disabled in this way, the message:

```
"Accounting suspended"
```

is printed on the system console.

acctsuspend is only relevant to systems which turn on process accounting.

Cost

None.

Dependencies (interactions with other system values)

`acctsuspend < acctresume`

argdevnblk

Name

`argdevnblk`—limit the size of the area reserved for `argdev`

Range

0	If 0, the area is configured dynamically.
48 to 256	Number of disk blocks (in multiples of 12). Each disk block equals 1 Kbyte.

Default

0

Use

`exec` uses `argdevnblk` as a scratch pad. The bigger the `argdevnblk`, the greater the number of concurrent `execs` allowed. Each `exec` call takes approximately 12 disk blocks.

Cost

Since `argdev` occupies the first part of the swap device, the greater the `argdevnblk`, the smaller the swap space actually available to the user's process.

Dependencies (interactions with other system values)

None.

dos_mem_byte

Name

`dos_mem_byte`—Reserves memory for the HP 98286 DOS Coprocessor

Range

0 to memory limited

Default

0

Use

The HP 98286 DOS Coprocessor uses system memory. This memory must be reserved on Series 300 systems with more than 6 Mbytes of main memory. The reserved memory is not available to HP-UX regardless of whether the DOS Coprocessor is running.

On Series 300 with 6 Mbytes or less, reserved memory is not required. The DOS Coprocessor shares system memory with HP-UX.

The amount of reserved memory depends upon the memory configuration of the DOS Coprocessor. 1 Mbyte of reserved memory is sufficient for DOS configured with 640 Kbytes of main memory and 64 Kbytes of expanded memory. Use of more DOS expanded or DOS extended memory requires correspondingly more reserved memory.

Cost

The amount of memory specified is reserved only for use by the DOS Coprocessor and is unavailable for other use.

Dependencies

none

dskless_cbufs

Name

dskless_cbufs—number of pages allocated to the cbuf (cluster buffer) pool

Range

0 to 128

Default

0	If the cluster code is not configured into the kernel.
dskless_mbufs*2+1	If the cluster code is configured into the kernel. Unless you change other configurable parameters, this value will be 12.

The cluster code is configured into the kernel if the kernel description file contains the driver called dskless. Refer to the section on kernel configuration for more information.

Use

This parameter determines the number of pages allocated to the cnode's **cbuf** pool. The cbuf pool, along with the cluster mbuf pool, are used for both inbound and outbound traffic.

On the root server, the cbuf pool is allocated when the cluster command is issued. On a cluster cnode, the cbuf pool is allocated at kernel initialization time.

Cost

Each cbuf is 1 Kbyte in length. There are four cbufs per page since each page on a Series 300 is 4 Kbytes. The dskless_cbufs parameter is actually a multiplier used as follows:

```
"dskless_cbufs" × "pagesize"
```

Dependencies (interactions with other system values)

None. The cluster pool is only utilized by the cluster code. The size of the cluster `cbuf` pool is directly dependent upon the `dskless_mbufs` parameter (`dskless_mbufs*2+1`).

dskless_fsbufs

Name

dskless_fsbufs—size of the file system buffers that can be allocated under interrupt

Range

0 to 200

Default

0	If the cluster code is not configured into the kernel.
serving_array_size	If the cluster code is configured into the kernel.

The cluster code is configured into the kernel if the kernel description file you used contains the driver called dskless. Refer to the section on kernel configuration for more information.

Use

This parameter directly sizes the cluster fsbuf page pool. The fsbuf pool is a collection of file system buffers that are utilized for inbound cluster traffic. The value should be <= serving_array_size. If you configure it to be > serving_array_size, it will be reset to be equal to serving_array_size. Diskless cnodes receive much less traffic than the root server, so this parameter should be smaller on the cluster cnode than on the root server.

If an inbound cluster message is large enough to require one of these buffers and none are available, the cluster protocol correctly handles retries.

On the root server the fsbuf pool is allocated when the root server issues the cluster command. On a cluster cnode, the netbuf pool is allocated at kernel initialization time.

Cost

`pagesize` × `dskless_fsbufs`

`pagesize` on the Series 300 is 4 Kbytes. Pages used for `fsbufs` are unavailable for general use.

Dependencies (interactions with other system values)

By default, `dskless_fsbufs` equals `serving_array_size`. It should always be <= the `serving_array_size` parameter.

dskless_mbufs

Name

dskless_mbufs—number pages allocated to the mbuf pool

Range

0 to 64

Default

0	If the cluster code is not configured into the kernel.
((server_node* num_cnodes)/4)+1	If the cluster code is configured into the kernel. Unless you change other configurable parameters, this value will be 6.

Note that the cluster code is configured into the kernel if the kernel description file contains the driver called dskless. Refer to the section on kernel configuration for more information.

Use

This parameter determines the number of pages allocated to the cluster mbuf. The term mbuf is taken from the Berkeley system, but this implementation is used specifically for HP's cluster cluster, and does not directly correspond to the Berkeley mbuf. The mbuf pool is a collection of small buffers used for both inbound and outbound cluster traffic.

On the root server, the mbufs are allocated when the root server issues the cluster command. On a cluster cnode, the mbuf pool is allocated at kernel initialization time.

Cost

`pagesize` × `dskless_mbufs`

On the Series 300, `pagesize` is 4 Kbytes.

Since these pages are allocated for cluster traffic, they are unavailable for general use.

Dependencies (interactions with other system values)

The cluster mbuf pool size is dependent on the parameters `server_node` and `num_cnodes`. The cluster `mbuf` pool is only utilized by the cluster code.

dskless_node

Name

`dskless_node`—flag used to let system know if cluster cnode or root server.

Range

1	If 1, this is a cluster cnode.
0	If 0, this is a server node.

Default

0

Use

If you are configuring the root server node, `dskless_node` should be 0 (default) and `server_node` should be set to 1. If you are configuring a standalone system, both `dskless_node` and `server_node` should be 0 (the default).

Cost

None. There is no direct cost for this flag.

Dependencies (interactions with other system values)

None. This parameter is not used; it is reserved for use in future releases.

dst

Name

dst—daylight savings time

Range

0 to 5

Default

1

Use

This variable defines the daylight savings time correction:

```
(from file ../h/time.h)

#define DST_NONE      0      /* not on dst */
#define DST_USA       1      /* USA style dst */
#define DST_AUST      2      /* Australian style dst */
#define DST_WET       3      /* Western European dst */
#define DST_MET       4      /* Middle European dst */
#define DST_EET       5      /* Eastern European dst */
```

The `dst` and `timezone` parameters are used in conjunction with the `gettimeofday` system call. They do not influence the `date` command.

Cost

None.

Dependencies (interactions with other system values)

It is used in conjunction with time zone.

filesizelimit

Name

`filesizelimit`—sets the file size limit for process 0

Range

0x00000010 to resource limited

Default

`0x1ffffff`

Use

This parameter sets the file size limit for process 0. The limit is given in units of 512-byte blocks and is inherited by child processes. Setting this parameter, effectively sets both the maximum and current filesize for all processes. If this parameter is not set by the user when configuring a kernel, the default setting will be `0x1ffffff`, which was the value it had before the parameter was tunable. After boot-up, the file size limit can always be changed for a process and its children, using the system call, `ulimit`. (Refer to the *ulimit(2)* entry in the *HP-UX Reference*.)

Cost

None.

Dependencies (interactions with other system values)

None.

fpa

Name

`fpa`—HP 98248A Floating Point Accelerator enable/disable

Range

0,1

Default

1

Use

`fpa` determines whether the code for the HP 98248A Floating Point Accelerator is included in the kernel. If `fpa` = 1 the the code is included, if `fpa` = 0 it is not.

If you wish to use the floating point capabilities, you must:

- set `fpa` to 1 (the default)
- create a device file for the card (described in the HP 98248A installation note)
- download the floating point microcode (described in the HP 98248A installation note)

Cost

Approximately 5 Kbytes.

Dependencies (interactions with other system values)

None.

maxdsiz

Name

`maxdsiz`—maximum data size

Range

262 144 bytes to 31 457 280 bytes (256 Kbytes to 30 Mbytes)

To increase beyond 30 Mbytes, see Table A-1 at the end of this chapter.

Default

0x01000000 (16 777 216 bytes, or 16 Mbytes)

Use

This value is entered in bytes.

`maxdsiz` defines the maximum size that the data segment of an executing process can have.

The default value is large enough for the data used by most processes. `maxdsiz` should only be increased if you know that you have one or more processes which use huge amounts of data.

Each time the system loads a process, or an executing process attempts to expand its data segment, the system checks the size of the process's data segment. If the `maxdsiz` is exceeded, the process will be terminated or returned with an appropriate error.

Cost

None.

Dependencies (interactions with other system values)

process text + process data + process stack <= 16 Mbytes (Model 310)

maxssiz

Name

`maxssiz`—maximum stack size

Range

262 144 bytes to 31 457 280 bytes (256 Kbytes to 30 Mbytes)

To increase beyond 30 Mbytes, see Table A-1 at the end of this chapter.

Default

0x00200000 (2 097 152 bytes, or 2 Mbytes)

Use

This value is entered in bytes.

`maxssiz` defines the maximum size that the stack segment of an executing process can have.

The default is large enough for the stack of most processes. `maxssiz` should only be increased if you have one or more processes which require a huge stack.

The stack grows dynamically. As the stack grows, the system checks the size of the process's stack segment. If the `maxssiz` is exceeded the process will be killed.

Cost

None.

Dependencies (interactions with other system values)

process text + process data + process stack <= 16 Mbytes (Model 310)

maxswapchunks

Name

`maxswapchunks`—limit the amount of swap on a workstation

Range

1 to infinity

Default

512

Use

This parameter limits the amount of swap space an individual workstation can have. This limiter should be used for a cluster cnode that tends to use a lot of the swap space connected to the root server. There is no reason to limit the amount of swap space if the cluster cnode has its own swap devices.

This limiter should be set to be larger than total swap space on a standalone system, or on a cluster cnode with local swap space. The maximum amount of swap space for the workstation is **maxswapchunks** × **dmmax** × 1024 bytes. Using default values of **maxswapchunks** and **dmmax** gives 256 Mbytes of swap space.

It is possible to get chunks of swap space at the end of the disk that are less than **dmmax** in size. Therefore, the minimum amount of swap space that a workstation can have is:

```
(maxswapchunks - N) × dmmax
```

where N is the number of swap devices.

Cost

$6 \times$ `maxswapchunks` bytes

Dependencies

The total swap space is limited by `maxswapchunks` $\times$ `dmmax` $\times$ 1024 bytes

maxtsiz

Name

`maxtsiz`—maximum text size

Range

262 144 bytes to 20 971 520 bytes (256 Kbytes to 20 Mbytes)

To increase beyond 20 Mbytes, see Table A-1 at the end of this chapter.

Default

0x01000000 (16 777 216 bytes, or 16 Mbytes)

Use

This value is entered in bytes.

`maxtsiz` defines the maximum size that the shared text segment of an executing process can have.

The current default will accommodate the text segments of most processes. Unless you run a process with a text segment larger than 16 Megabytes, `maxtsiz` should not be modified.

Each time the system loads a process with shared text, the system checks the size of its shared text segment. `exec` issues an error message and aborts the process if the process's text segment is larger than `maxtsiz`.

Cost

None.

Dependencies (interactions with other system values)

process text + process data + process stack <= 16 Mbytes (Model 310)

maxuprc

Name

`maxuprc`—maximum number of user processes

Range

3 to (nproc - 3)

Default

50 processes

Use

`maxuprc` defines the maximum number of simultaneous processes that a user can have. A user is identified by the UID (user ID) number, not by the number of login instances. Each user will need at least one process for the shell, plus an adequate number to be able to do useful work (25 is usually more than enough).

The superuser is exempt from this limit.

Pipelines need at least one simultaneous process for each side of a '|'. Some commands, such as `cc`, `fc`, and `pc`, use more than one process per invocation.

When the total number of processes for a given user is larger than `maxuprc`, the following message is printed to the user tty:

`no more processes`	*if from shell, or*
`EAGAIN error`	*if a fork call from a program.*

Cost

none.

Dependencies (interactions with other system values)

If `maxuprc` is set to a value greater than or equal to `nproc` (maximum number of processes in the system) then `maxuprc` is no longer a limit, and the system can be hoarded by a single user.

maxusers

Name

maxusers—limiter for system resource allocation

Range

0 to memory limited

Default

8

Use

maxusers is a limiter for system resource allocation. The global operating system parameters that depend on maxusers are listed in the "Dependencies" section. maxusers is *not* a limiter for the number of users in the system, therefore its name (inherited from System V) is confusing.

Cost

By itself, maxusers does not determine the size of any structures in the system. However, the value of maxusers indirectly affects kernel data structures. Consequently, the amount of memory available for user processes will be smaller with higher values of maxusers. Data structures that depend on the value of maxusers will grow as follows. (Although "maxusers" affects the "nclist" variable, "nclist" is not considered a tunable parameter. All other variables listed here are tunable and each is described elsewhere in this chapter.)

```
32 bytes × nclist
180 bytes × nproc
16 bytes × ncallout
286 bytes × ninode
30 bytes × nfile
188 bytes × ntext
20 bytes × serving_array_size
```

Dependencies (interactions with other system values)

nclist = 100+(16×maxusers)

nproc = 20 + (8 × maxusers) + ngcsp

ncallout = 16 + nproc + using_array+size + serving_array_size

ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)

nfile = 16 × ((nproc+16+maxusers) ÷ 10) +32+ (2 × npty)

ntext = 40 + maxusers

serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)

mesg

Name

`mesg`—System V messages

Range

0 to 1

Default

1

Use

`mesg` determines whether the code for System V IPC messages is included in the kernel. If `mesg`=1, the code is included; if `mesg`=0, then the code is not included. If `mesg`=0, all programs using the system calls (refer to the entries for *msgget(2), msgbp(2),* and *msgctl(2)* in the *HP-UX Reference*) will receive a SIGSYS signal.

Cost

2680 bytes.

Dependencies (interactions with other system values)

All message parameters depend on mesg.

If `mesg` = 0, then the message parameters are not declared and no message structures are allocated.

If `mesg` = 1, then the message parameters are declared and can can be tuned and all other message structures are allocated.

minswapchunks

Name

`minswapchunks`—set the minimum amount of swap space on a system

Range

1 to infinity

Default

4

Use

This parameter is used to set the minimum amount of swap space for a system. The minimum amount of swap space would be:

```
minswapchunks × dmmax × 1024 bytes
```

Cost

There is no cost for this parameter.

Dependencies (interactions with other system values)

Must not be greater than `maxswchunks`. Set to `maxswchunks` if larger.

Total swap space must be at least `minswapchunks` × `dmmax`

msgmap

Name

msgmap—message map

Range

3 to memory limited

Default

(msgtql+2)

Use

Each set of messages allocated per identifier occupies one or more contiguous slots in the msg array. As messages are allocated and deallocated the msg array may become fragmented.

`msgmap` dimensions the resource map used to allocate the buffer space for messages. This map shows the free holes in the `msg` array. An entry in the map is used to point to each set of contiguous unallocated slots, and it consists of a pointer to the set, plus the size of the set.

If message usage is heavy, and a request for a message set cannot be accommodated, the message:

```
danger: mfree map overflow
```

will appear. If you get the error, you should make a new kernel with a larger `msgmap`.

There will be less fragmentation of the `msg` array if all message identifiers have the same number of messages. `msgmap` can then be smaller.

Three (3) is the lower limit: one slot is overhead for the map and the second slot is always needed at system initialization to indicate that the `msg` array is completely free.

Cost

4 bytes×`msgmap`

Dependencies (interactions with other system values)

`msgmap` <= (`msgtql`+2)

`msgmap` <= (`msgseg`+2)

If `mesg` = 0, then the message code is not included in the kernel and the value of `msgmap` is irrelevant.

If `mesg` = 1, then the message code is included and `msgmap` can then be tuned.

msgmax

Name

msgmax—message maximum size

Range

0 to memory limited

Default

8192 bytes

Use

msgmax limits the size, in bytes, of a single message.

It should be increased only if you have applications which require larger messages. Its main value is to keep malicious or poorly written programs from using all the message buffer space.

A msgsnd system call which attempts to send a message larger than msgmax bytes returns the EINVAL error.

Cost

None.

Dependencies (interactions with other system values)

msgmax <= msgmnb

If mesg = 0, then the message code is not included in the kernel and the value of msgmax is irrelevant.

If mesg = 1, then the message code is included and msgmax can then be tuned.

msgmnb

Name

msgmnb—maximum number of bytes on message queue

Range

0 to memory limited

Default

16 384 bytes

Use

msgmnb is the maximum total size, in bytes, of all messages that can be queued on a message queue at the same time.

A msgsnd system call which attempts to exceed this limit either:

- blocks or,
- returns EAGAIN error

depending on whether the IPC_NOWAIT flag is set with the call.

Cost

None.

Dependencies (interactions with other system values)

msgmnb >= msgmax

msgmnb <= (msgssz×msgseg)

If mesg = 0, then the message code is not included in the kernel and the value of msgmnb is irrelevant.

If mesg = 1, then the message code is included and msgmnb can then be tuned.

msgmni

Name

`msgmni`—number of message queue identifiers

Range

1 to memory limited

Default

50

Use

`msgmni` dimensions the array of message queue identifiers.

A message queue identifier is needed for each message queue in the system.

When `msgmni` message queues already exist, an attempt to allocate a new message queue with the `msgget` system call returns an ENOSPC error.

If a message queue is not deallocated, it will continue to exist in the system even after the process(es) using it have died.

When you are through using a message queue, remember to deallocate them. Running into the `msgmni` limit usually means that users have not freed them up.

Cost

46 bytes×`msgmni`

Dependencies (interactions with other system values)

If `mesg` = 0, then the message code is not included in the kernel and the value of `msgmni` is irrelevant.

If `mesg` = 1, then the message code is included and `msgmni` can then be tuned.

msgseg

Name

`msgseg`—message segments

Range

1 to 32 767

Default

16 384

Use

`msgseg` together with `msgssz` determines the size of the buffer available for queuing messages.

`msgssz` determines the size, in bytes, of the units in which messages are allocated space. When a message is allocated, its size is rounded up to the nearest multiple of `msgssz`.

`msgseg` is the number of these units available.

In most cases the product of `msgseg`×`msgssz` is of most interest since it determines the total amount of space available for messages. Different `msgseg:msgssz` ratios which yield the same product will just cause this space to be fragmented differently for the same usage.

Cost

`msgssz` bytes×`msgseg`

Dependencies (interactions with other system values)

$(\texttt{msgseg} \times \texttt{msgssz}) >= \texttt{msgmnb}$

$(\texttt{msgseg} \times \texttt{msgssz}) >= \texttt{msgmax}$

If `mesg` = 0, then the message code is not included in the kernel and the value of `msgseg` is irrelevant.

If `mesg` = 1, then the message code is included and `msgseg` can then be tuned.

msgssz

Name

msgssz—message segment size

Range

1 to memory limited

Default

1 byte

Use

msgssz together with msgseg determines the size of the buffer available for queueing messages.

msgssz determines the size, in bytes, of the units in which messages are allocated space. When a message is allocated, its size is rounded up to the nearest multiple of msgssz.

msgseg is the number of these units available.

In most cases, the product of msgseg×msgssz is of most interest since it determines the total amount of space available for messages. Different msgseg:msgssz ratios which yield the same product will just cause this space to be fragmented differently for the same usage.

Cost

msgseg×msgssz bytes

Dependencies (interactions with other system values)

(msgseg×msgssz) >= msgmnb

(msgseg×msgssz) >= msgmax

If mesg = 0, then the message code is not included in the kernel and the value of msgssz is irrelevant.

If mesg = 1, then the message code is included and msgssz can then be tuned.

msgtql

Name

msgtql—number of message headers

Range

1 to memory limited

Default

40

Use

msgtql dimensions an array of message headers. A message header is used for each message queued in the system.

A msgsnd system call which attempts to exceed this limit either:

- blocks or,
- returns EAGAIN error

depending on whether the IPC_NOWAIT flag is set with the call.

Cost

12 bytes×msgtql

Dependencies (interactions with other system values)

msgmap <= msgtql + 2

If mesg = 0, then the message code is not included in the kernel and the value of msgtql is irrelevant.

If mesg = 1, then the message code is included and msgtql can then be tuned.

nbuf

Name

`nbuf`—number of buffer headers

Range

0	At boot time, if `nbuf` = 0, the amount of buffer space is dynamically set to 10 percent of available memory for the first 5 Mbytes, 5 percent for the remaining memory.
16 to 768	The number of file system buffer headers is set to the specified value.

Default

0 (configured dynamically)

Use

The file system buffer cache is a pool of pages (bufpages) and a configurable number of buffer headers. `nbuf` defines the number of file system buffer headers.

`nbuf` can be left alone, in which case the default values will be used. However, `nbuf` can be used to decrease or increase the buffer cache within the ranges given above (refer to Range). These two cases are discussed below:

If you have a Model 310, it assigns (in pages) to the system buffer cache 10% of memory for the first 2 Mbytes. If the amount of RAM is greater than 2 Mbytes, use an `nbuf` of 51, which is the default for 2 Mbytes.

- The default case:

  At boot time, as the system dynamically allocates memory, it assigns (in pages) to the system buffer cache 10 percent of the available memory for the first 5 Mbytes, 5 percent for the remaining memory[1]. Then, `nbuf` is computed so that each `nbuf` header has one 4 Kbyte (4096 bytes) bufpages of space assigned to it.

For example, on a Series 300 with 8 Mbytes of RAM, the default is nbuf = 0. After the dynamic allocation, nbuf = 166 and the buffer cache = 166×4096 bytes.

- To decrease or increase the buffer cache:

  The user can specify the size of the buffer cache, independent of memory size. For instance, to specify 127 buffer headers, the sample dfile entry would be:

  ```
  * tunable system parameters
     nbuf     127
  ```

  Notice that in the above example, the size (127) may be smaller or larger than the default case.

These buffers are used for all file system I/O operations, plus all other block I/O operations in the system (for example, exec, mount, i-node reading, some device drivers, etc.).

Cost

4096 bytes×each bufpage

104 bytes×each nbuf header structure

104 bytes×max(256, (nbuf÷2))

Dependencies (interactions with other system values)

None.

ncallout

Name

`ncallout`—number of callouts

Range

6 to memory limited

Default

16 + `nproc` + `using_array_size` + `serving_array_size`

Use

`ncallout` is the maximum number of timeouts which can be scheduled by the kernel at any one time. Timeouts are used by :

- "alarm" (system call)
- `setitimer` (system call)
- drivers
- `uucp` processes
- process scheduling
- HP-UX cluster kernel code

When the system runs out of callouts, it prints the following fatal error to the console:

```
panic: timeout table overflow
```

Cost

16 bytes×`ncallout`

Dependencies (interactions with other system values)

The larger `nproc` is, the larger `ncallout` should be.

A rough guideline of 1 callout per process should be used unless one has processes which use lots of callouts.

ndilbuffers

Name

`ndilbuffers`—number of DIL buffers

Range

1 to memory limited

Default

30

Use

`ndilbuffers` defines the maximum number of DIL open device files at any one time in the system.

`ndilbuffers` is used exclusively by the Device I/O Library. If DIL is not used, no DIL buffers are necessary.

Cost

314 bytes×`ndilbuffers`

Dependencies (interactions with other system values)

None.

netmeminit

Name

`netmeminit`—number of bytes of memory to be preallocated at system initialization time

Range

0 to 512 × `NETCLBYTES` (1000 Kbytes)

Default

0

Use

The `netmeminit` parameter specifies the number of bytes of memory to be preallocated for networking at system initialization time. There is no delay associated with networking getting physical memory from the system.

Dependencies (interactions with other system values)

`netmeminit < netmemmax`

netmemmax

Name

`netmemmax`—network memory pool maximum

Range

75 000 to 1 000 000 bytes

Default

250 000 bytes

Use

`netmemmax` defines the maximum network memory pool size. The network memory pool starts with minimal memory and grows as the amount of network activity grows, until `netmemmax` is reached. If more network memory is needed after the limit is reached, the network operation requesting the memory receives an "out-of-memory" message and the operation fails.

Cost

The network memory pool is allocated from the available system memory. Larger values of `netmemmax` allow the network to acquire more memory, leaving less memory available for other system activities.

Dependencies (interactions with other system values)

If `netmemmax` is less than or equal to `netmemthresh`, network memory reservation is always in effect.

netmemthresh

Name

`netmemthresh`—network memory reservation threshold

Range

-1 to 2 147 483 647

Default

100 000

Use

`netmemthresh` defines how large the network memory pool is allowed to grow before disabling network memory reservation.

As long as network memory reservation is in effect (i.e., `netmemthresh` has not been reached), each connection is guaranteed its maximum memory requirements. When the cumulative effect of each connection's memory reservations meets or exceeds `netmemthresh`, network memory reservation is "disabled". Subsequently, all connections are relegated to memory allocation on an "as needed" basis. Because connections are no longer guaranteed memory allocation, connections could be lost due to insufficient memory.

Network memory reservation can be enforced regardless of the size of the network memory pool by assigning `netmemthresh` the value -1.

Cost

None.

Dependencies (interactions with other system values)

If `netmemthresh` is less than `netmemmax`, the network memory pool typically grows very quickly to `netmemthresh` size and remains there, except under extremely heavy network loads. Increasing the value of `netmemthresh` reserves more memory under light network loads, leaving less memory available to other system activities.

If `netmemthresh` is greater than or equal to `netmemmax`, or `netmemthresh` equals -1, network memory reservation is in effect at all times. In this case, the network memory pool typically grows to `netmemmax` size very quickly. Memory requirements of all connections are guaranteed, but fewer connections can be supported than with lower values of `netmemthresh`.

nfile

Name

`nfile`—number of files

Range

14 to memory limited

Default

16 × ((nproc+16+maxusers) ÷ 10) +32+ (2 × npty)

Use

`nfile` defines the maximum number of open files at any one time in the system.

It is the number of slots in the file descriptor table. Be generous with this number since the cost is low, and not having enough slots would cut down in the amount of work that can be done simultaneously in the system.

Cost

30 bytes×`nfile`

Dependencies (interactions with other system values)

Windows/9000	3 or 4 file descriptors per window
Processes	At least 3 file descriptors per process (stdin, stdout, stderr)
Pipes	2 per pipe (1 per side)

Shell scripts and background processes are treated like any other processes.

`nfile` depends on `nproc`, `maxusers`, and `npty` as shown in the "Default" section.

nflocks

Name

nflocks—number of file locks

Range

2 to 2000

Default

200

Use

nflocks gives the possible number of file/record locks in the system. When choosing this number it should be noted that one file may have several locks and that especially databases may need a large number of locks (if they use lockf at all).

Cost

22 bytes×nflocks

Dependencies (interactions with other system values)

None.

ngcsp

Name

ngcsp—number of general cluster server processes (for cluster server)

Range

0 to memory limited

Default

8 × num_cnodes

The default for num_cnodes is 0, so unless you have changed num_cnodes the value for ngcsp is also 0.

Use

ngcsp specifies the maximum number of general cluster server processes (GCSPs) that can exist simultaneously in the system. GCSPs are required only on the cluster's root server, so the value of ngcsp should be zero on machines that run standalone or cluster.

GCSPs are started with the csp command; the command returns an error if asked to start more than ngcsp processes. The system will sometimes attempt to spawn more GCSPs than requested by this command to keep up with heavy demands. Thus it is recommended, but not required, that the value of ngcsp be about twice the value specified for the csp command.

Cost

18 bytes × ngcsp

Dependencies (interactions with other system values)

ngcsp $= 8 \times$ num_cnodes

nproc $= 20 + 8 \times$ maxusers $+$ ngcsp

Each GCSP requires one process slot. Therefore, ngcsp should be less than (nproc - 5).

Each GCSP in use requires a serving array entry as do some other requests, therefore ngcsp must be less than serving_array_size.

ninode

Name

`ninode`—number of inodes

Range

14 to memory limited

Default

`((nproc+16+maxusers)+32+(2×npty)+server_node×18×num_cnodes)`

Use

`ninode` defines the maximum number of open inodes which can be in-core.

`ninode` specifies the maximum number of slots in the in-core inode table. This in-core table is used for fast access of inode information. For efficiency reasons, the table is hashed.

Cost

286 bytes×`ninode`

Dependencies (interactions with other system values):

Each unique open file has an open inode associated with it, therefore the larger the number of unique open files, the larger `ninode` should be. If this is the root server of an HP-UX cluster this pertains to all opens on the entire cluster.

The `ninode` parameter depends on `nproc`, `maxusers`, `npty`, `server_node` and `num_cnodes` as shown in the equation in the "default" section.

nproc

Name

nproc—number of processes

Range

6 to memory limited

Default

20+8×maxusers + ngcsp

maxusers defaults to 8 and ngcsp defaults to 0, so this value defaults to 84.

Use

nproc specifies the maximum total number of processes that can exist simultaneously in the system.

There have to be at least three (3) system overhead processes at all times (swapper, init, and pageout demon), and there is always one entry reserved for the superuser.

Examples of other features and the number of processes are:

LAN network	3 processes
Printer scheduler	1 process
Cron	1 process

Your application's manual may discuss the number of processes used. For a discussion of the number of processes used by HP Windows/9000 refer to the *HP Windows/9000 User's Manual,* the appendix "Window Limitations".

When the total number of processes in the system is larger than nproc, the following messages are printed:

- To the console:

```
proc: table is full
```

- To the tty of the user who tries to start the last process(es):

`no more processes`	*if from shell, or*
`EAGAIN error`	*if a fork call from a program.*

Cost

180 bytes×nproc

Dependencies (interactions with other system values)

```
ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)

nfile = 16 × ((nproc+16+maxusers) ÷ 10) +32+ (2 × npty)

ncallout = 16 + nproc + using_array_size + serving_array_size

maxuprc <= (nproc - 3)

nproc = 20 + (8 × maxusers) + ngcsp
```

using_array_size is directly dependent upon the nproc configurable parameter

npty

Name

`npty`—number of pseudo-teletypes

Range

1 to memory limited

Default

82

Use

`npty` controls the maximum number of pseudo-teletype (pty) devices in the system.

`npty` is used primarily by the window subsystem. If you use windows, you are a large user of `npty`. If you use **`remsh`**, **`vt`**, or **`rlogin`**, you are a small user of `npty`. If windows are not used, and the pty driver is not explicitly used, no pty structures are necessary.

Cost

188 bytes×`npty`

Dependencies (interactions with other system values)

None.

ntext

Name

ntext—number of shared text descriptors

Range

10 to memory limited

Default

40 + maxusers

Use

ntext defines the maximum number of shared text (or code) descriptors (data structures describing the text) which can be active at any one time. Note that this does not limit the number of processes sharing the same shared text program.

Attempting to start a new process that may require a new text descriptor when no more text descriptors are available will get "text table is full" message on the console and the new process will be killed.

Cost

192 bytes×ntext

Dependencies (interactions with other system values)

The default for ntext is 40 + maxusers.

num_cnodes

Name

`num_cnodes`—limiter for cluster system resource allocation

Range

0 to 255

Default

0

Use

`num_cnodes` is used as a limiter for cluster resource allocation; the value of other global system parameters depends on `num_cnodes`. By itself it does not determine the size of any structures in the system.

It is used as an indicator of the number of cluster cnodes that a server can reasonably expect to serve simultaneously. It is *not* used as a limit for the number of cnodes supported by the server.

Suggested values are:

```
5 if  0- 5 cnodes
10 if  6-10 cnodes
15 if 11-15 cnodes
20 if 16-20 cnodes
```

Cost

There is no direct cost for this flag.

A higher value of `num_cnodes` indirectly results in larger kernel data structures and buffers; consequently, the amount of memory available for user processes becomes smaller.

Dependencies (interactions with other system values)

serving_array_size = (server_node $\times$ num_cnodes $\times$ maxusers + 2 $\times$ maxusers)

ninode = (nproc + 16 + maxusers)+32+(2 $\times$ npty)+(server_node $\times$ 18 $\times$ num_cnodes)

ngcsp = 8 $\times$ num_cnodes

dskless_mbufs = (((server_node $\times$ num_cnodes)/4)+1)

num_lan_cards

Name

`num_lan_cards`—maximum number of LAN interface cards the system will support

Range

0 to 5

Default

2

Use

`num_lan_cards` defines the maximum number of LAN interface cards the system will support. This number should be greater than or equal to the number of LAN interface cards actually present.

Cost

Approximately 1200 bytes per LAN interface card.

Dependencies (interactions with other system values)

None.

parity_option

Name

`parity_option`—used to handle parity errors

Range

0 to 2

Default

2

Use

`parity_option` selects the kind of action that the system takes if it encounters a parity error.

The actions are as follows:

0 Print a 'Parity error' message to the console.

1 Print a 'Parity error' message to console, plus:

- if user state, it kills the current process (which may not always be the process which caused the error, as with a DMA card) and prints an error message to its tty.
- if supervisor state, it panics with a 'parity error' message to the console.

2 Always panics with a 'parity error' message to console.

Caution Values other than 2 could result in data corruption depending on where the RAM parity error occurs.

Cost

None.

Dependencies (interactions with other system values):

None.

scroll_lines

Name

`scroll_lines`—total number of ITE text lines

Range

100 to 5000

Default

100

Use

The ITE (Internal Terminal Emulator) will remember more lines of text than can be displayed on the screen. This parameter controls how many lines of text will be remembered.

For example, on a 46 line screen, if `scroll_lines` is 460, then there will be a total of ten screens of information available, of which only one is displayed at a time.

Cost

2 × <screen width in characters> × `scroll_lines`

Dependencies (interactions with other system values)

None.

selftest_period

Name

`selftest_period`—period between execution of kernel selftest

Range

0, 90 to 3600 seconds

Default

0	If the cluster code is not configured into the kernel.
120 sec.	If the cluster code is configured into the kernel.

Note that the cluster code is configured into the kernel if the kernel description file contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

Use

The cluster kernel's selftest code uses the `selftest_period` parameter to determine how often to run a selftest. The selftest checks the availability of kernel resources required for cluster operation. The parameter is interpreted as the minimum number of seconds between execution of the selftest check. A 0 value indicates that the selftest should not be run.

Cost

None.

Dependencies (interactions with other system values)

None.

sema

Name

sema —System V semaphores

Range

0 to 1

Default

1

Use

sema determines whether the code for System V IPC semaphores will be included in the kernel.

If sema = 1, the code is included;
if sema = 0, then the code is not included.

HP Windows/9000 and the Starbase graphics library both require the semaphore code.

If sema=0, and you have programs which use the *semget(2)* or *semop(2)* system calls, you will receive a SIGSYS signal.

Cost

4240 bytes.

Dependencies (interactions with other system values)

All semaphore parameters depend on the value of sema.

If sema = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

If sema = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

Note If `sema` = 0, then HP Windows/9000 and Starbase graphics will be unable to run.

semaem

Name

semaem—"adjust on exit" maximum value for semaphores

Range

0 to min (semvmx, 32767)

Default

16 384

Use

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

semaem is the maximum value by which a semaphore can be undone.

This value is cumulative per process, so if one process has more than one undo operation on a semaphore, the value of each undo operation is added up in the variable semadj. semadj is the number by which the semaphore will be incremented or decremented if the process dies.

Refer to the entry for *semop(2)* in the *HP-UX Reference* for more detailed information on semaphore undos.

Any semop calls which attempt to set |semadj| > semaem result in an ERANGE error.

Cost

None.

Dependencies (interactions with other system values)

$|$`semaem`$|$ <= `semvmx`

If `sema` = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema` = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

semmap

Name

`semmap`—semaphore map

Range

4 to memory limited

Default

semmni + 2

Use

Each set of semaphores allocated per identifier occupies one or more contiguous slots in the `sem` array. As semaphores are allocated and deallocated the `sem` array may become fragmented.

`semmap` dimensions the resource map which shows the free holes in the `sem` array. An entry in this map is used to point to each set of contiguous unallocated slots, and it consists of a pointer to the set, plus the size of the set.

If semaphore usage is heavy, and a request for a semaphore set cannot be accommodated, the message:

```
danger: mfree map overflow
```

will appear. It will then be helpful to make a new kernel with a larger `semmap`.

Fragmentation of the `sem` array is reduced if in usage of the system all semaphore identifiers have the same number of semaphores. `semmap` can then be somewhat smaller.

Four (4) is the lower limit: 1 slot is overhead for the map and the second slot is always needed at system initialization to indicate that the sem array is completely free.

Cost

4 bytes×`semmap`

Dependencies (interactions with other system values)

(`semmap`-2) = maximum number of contiguous unallocated pieces of the `sem array`.

`semmap` <= (`semmni`-2)

If `sema` = 0, then the semaphore code is not included in the kernel and the value of `semmap` is irrelevant.

If `sema` = 1, then the semaphore code is included and `semmap` can then be tuned.

semmni

Name

semmni—semaphore number of identifiers

Range

2 to memory limited

Default

64

Use

semmni defines the number of sets (identifiers) of semaphores available to the users.

When the system runs out of semaphore sets, the semget system call will return a ENOSPC error message.

Cost

30 bytes$\times$semmni

Dependencies (interactions with other system values)

semmni <= semmns

semmns <= (semmni$\times$semmsl)

semmsl = 50 /* maximum number of semaphores per ID, not tunable */

semmap <= (semmni+2)

If sema = 0, then the semaphore code is not included in the kernel and the value of semmni is irrelevant.

If sema = 1, then the semaphore code is included and semmni can then be tuned.

semmns

Name

semmns—semaphore number in system

Range

2 to memory limited

Default

128

Use

semmns defines the total number of semaphores available to the users of the system.

When the system does not have enough contiguous semaphores in the sem array to satisfy a semget request, the call returns a ENOSPC error. This error may occur even though there may be enough free semaphores, but they are not contiguous.

Cost

8 bytes×semmns

Dependencies (interactions with other system values)

semmni <= semmns

semmns <= (semmni×semmsl)

semmsl = 50 /* maximum number of semaphores per ID, not tunable */

If sema = 0, the semaphore code is not included in the kernel and the value of semmns is irrelevant.

If sema =1, the semaphore code is included and then semmns can be tuned.

semmnu

Name

semmnu—semaphore number of undo structures

Range

1 to (nproc - 3)

Default

30

Use

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

`semmnu` is the number of processes which can have undo's pending on a given semaphore. It determines the size of the `sem_undo` structure.

Refer to the entry for *semop(2)* for a more detailed explanation of undos.

You should increase `semume` if the user gets ENOSPC errors on `semop` calls using the SEM_UNDO flag.

Cost

(6+(8×`semume`)) bytes×`semmnu`

Dependencies (interactions with other system values)

`semmnu` determines the size of the structure sem_undo, which in turn contains the substructure dimensioned by `semume`.

There is no point in having `semmnu` > (nproc-3) since it is the largest number of processes in the system which could use semaphores simultaneously.

If `sema` = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema` = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

semume

Name

semume—semaphore undo entries per process

Range

1 to semmns

Default

10

Use

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

semume limits the number of semaphores that each process can have undos pending on.

Refer to the entry for *semop(2)* for a more detailed explanation of undos.

When the user gets EINVAL errors on semop calls with the SEM_UNDO flag, then you should increase the value of semume.

Cost

(6+(8×semume)) bytes×semmnu

Dependencies (interactions with other system values)

semume <= semmns

semume is the size of the substructure undo, which is part of the sem_undo structure. The size of sem_undo is determined by semume.

If `sema` $= 0$, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema` $= 1$, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

semvmx

Name

`semvmx`—semaphore maximum value

Range

1 to 65 535

Default

32 767

Use

`semvmx` is the maximum value that a semaphore is allowed to reach. This limit is imposed by the largest number that can be stored in a 16-bit unsigned integer (65 535).

`semop` system calls which try to increment a semaphore's value > `semvmx` result in ERANGE errors. If semvmx > 65 535 then semaphore values can overflow and these errors will not be caught.

Cost

None.

Dependencies (interactions with other system values):

`semaem <= semvmx`

If `sema` = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema` = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

server_node

Name

server_node—flag used to help size an array for the root server's inbound requests.

Range

1 If 1, serving_array[] and ninode are sized for a server node.

0 If 0, serving_array[] and ninode are sized for a cluster cnode.

Default

0

Use

This parameter is used as a flag to determine the size of serving_array[] and ninode. serving_array[] is an array of kernel structures used by a cnode for inbound requests. If this parameter is set (server_node=1), serving_array[] will be sized appropriately for a root server node. The cluster's root server should have this parameter set to 1, and should have the dskless_node parameter set to 0 (default).

If you are configuring a cluster cnode, server_node should be 0 (default) and dskless_node should be set to 1. If you are configuring a standalone system, both dskless_node and server_node should be 0 (the default).

Refer to the serving_array_size and ninode operating system parameters for information on these resources.

Cost

None. There is not direct cost for this flag.

Dependencies (interactions with other system values)

Used in ninode and serving_array_size as follows:

```
serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)

ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)

dskless_mbufs = (((server_node × num_cnodes)/4)+1)
```

serving_array_size

Name

serving_array_size—size of the cluster's serving array

Range

0 to 200

Default

0	If the cluster code is not configured into the kernel.
(server_node*num_cnodes* maxusers)+(2*maxusers)	If the cluster code is configured into the kernel. Unless you change other configurable parameters, this value will be 16.

The cluster code is configured into the kernel if the kernel description file you used contains the driver called dskless. Refer to the section on kernel configuration for more information.

Use

serving_array_size defines the size of the kernel's serving array. serving_array[] is an array of kernel structures that holds information related to inbound requests. It is sized according to the following formula:

```
serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)
```

Each inbound request consumes a single serving_array[] entry that is held until the request is satisfied.

If both configurable parameters, dskless_node and server_node, are equal to zero, the system is treated as a standalone, non-clustered, kernel. In this situation, serving_array_size is zero and serving_array[] is not compiled into the kernel.

Cost

20 bytes × serving_array_size

Dependencies (interactions with other system values)

The size of the serving array is directly dependent upon the following parameters: server_node, num_cnodes, and maxusers.

serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)

By default, the number of file system buffers that can be allocated under interrupt is equal to the value of serving_array_size (see configurable parameter dskless_fsbufs).

Each GCSP requires a serving array entry (refer to ngcsp)

shmall

Name

`shmall`—total size of all shared memory segments

`shmall` is present just for compatibility reasons with System V.

shmbrk

Name

`shmbrk`—shared memory break

Range

0 to maximum number of pages in the data segment.

Default

16 pages

Use

`shmbrk` defines the size of the gap, in pages, between the top of the current data segment and the first default shared memory address.

Cost

None.

Dependencies (interactions with other system values)

If `shmem` = 0, then the code for shared memory is not included in the kernel and `shmbrk` has no effect.

If `shmem` = 1, then the code for shared memory is included in the kernel and `shmbrk` is tunable.

shmem

Name

`shmem`—System V shared memory enable/disable

Range

0, 1

Default

1

Use

`shmem` determines whether the code for System V IPC shared memory is included in the kernel. If `shmem` = 1 then the code is included, if `shmem` = 0 it is not.

Cost

Approximately 11 Kbytes.

Dependencies (interactions with other system values)

HP Windows/9000 and the Starbase graphics library both require the presence of shared memory code.

shmmax

Name

shmmax—shared memory maximum

Range

2 Mbytes to 6 Mbytes

To increase beyond 6 Mbytes

Default

0x00600000 (768 Mbytes)

Use

shmmax defines the maximum shared memory segment size.

Cost

None.

Dependencies (interactions with other system values)

shmmin <= shmmax

If shmem = 0, then the code for shared memory is not included in the kernel and the value of shmmax is irrelevant.

If shmem = 1, then the code for shared memory is included and shmmax is tunable.

For Windows/9000 to run, shmmax must be >= 2 Mbytes.

shmmaxaddr

Name

`shmmaxaddr`—shared memory maximum address

Range

positive integer

Default

0x01000000 (16 Mbytes)

Use

`shmmaxaddr` specifies the highest address allowed for a shared memory segment in user address space.

Cost

None.

Dependencies (interactions with other system values)

If `shmem` = 0, then the code for shared memory is not included in the kernel and `shmmaxaddr` has no effect.

If `shmem` = 1, then the code for shared memory is included and `shmmaxaddr` is tunable.

If you allow a shared memory segment to attach at a high address location, you will require more system overhead for the translation table data structure. *Do not reconfigure this unless absolutely necessary.*

shmmin

Name

shmmin—shared memory minimum

Range

positive integers

Default

1 byte

Use

shmmin defines the minimum shared memory segment size.

Cost

None.

Dependencies (interactions with other system values)

shmmin < shmmax

If shmem = 0, then the code for shared memory is not included in the kernel and the value of shmmin is irrelevant.

If shmem = 1, then the code for shared memory is included and shmmin is tunable.

If it is reconfigured other subsystems (such as Windows/9000) may not work.

shmmni

Name

shmmni—shared memory maximum number of identifiers

Range

positive integers

Default

30 identifiers

Use

shmmni defines the system-wide maximum number of shared memory segments.

shmmni should be large enough to hold as many shared memory segments as will be used simultaneously.

Cost

The data structure associated at each shared memory segment is about 100 bytes.

Dependencies (interactions with other system values)

If shmem = 0, then the code for shared memory is not included in the kernel and shmmni has no effect.

If shmem = 1, then the code for shared memory is included and shmmni can be tuned.

Windows/9000 and Starbase graphics require shmmni to be >= 4.

shmseg

Name

shmseg—shared memory segments

Range

positive integers

Default

10

Use

shmseg defines the maximum number of shared memory segments that can be attached to a process at any given time.

Cost

about 12 bytes×nproc×shmseg

Dependencies (interactions with other system values)

If shmem = 0, then the code for shared memory is not included in the kernel and shmseg has no effect.

If shmem = 1, then the code for shared memory is included and shmseg can be tuned.

Windows/9000 and Starbase graphics require that shmseg be >= 4.

timeslice

Name

`timeslice`—scheduling timeslice interval

Range

-1 to 2 147 483 647

Default

0

Use

The system performs round-robin scheduling among processes at a given priority. The timeslice interval is the amount of time one process is allowed to run before the CPU is given to the next process at the same priority.

The value of timeslice is specified in units of 20 millisecond clock ticks. There are two special values:

0	use the system default value (currently 5 ticks, or 100 milliseconds)
-1	disable round-robin scheduling completely

One side-effect of this parameter is that a process always checks for pending signals when its timeslice expires. This guarantees that a process which does not make any system calls (including a runaway process in an infinite loop) can be terminated. Thus setting timeslice to a very large value, or to -1, can prevent such processes from seeing signals.

It is anticipated that this parameter will only be changed on systems dedicated to applications with specific real-time needs.

Cost

There is no memory allocation related to this parameter. There is some amount of CPU time spent at each timeslice interval, but this time has not been precisely measured.

Dependencies (interactions with other system values)

None.

timezone

Name

`timezone`—minutes west of Greenwich

Range

-780 to 780

Default

420

Use

Time zone information indicates the minutes west of Greenwich.

```
struct  timezone tz = { TIMEZONE, DST };
struct timezone {
int tz_minuteswest; /* minutes west of Greenwich */
int tz_dsttime;     /* type of dst correction */
};
#define DST_NONE      0     /* not on dst */
#define DST_USA       1     /* USA style dst */
#define DST_AUST      2     /* Australian style dst */
#define DST_WET       3     /* Western European dst */
#define DST_MET       4     /* Middle European dst */
#define DST_EET       5     /* Eastern European dst */
```

The `dst` and `timezone` parameters are used in conjunction with the `gettimeofday` system call. They do not influence the `date` command.

Cost

None.

Dependencies (interactions with other system values)

It is used in conjunction with `dst` (daylight savings time).

unlockable_mem

Name

unlockable_mem—unlockable memory

Range

0 to the available memory indicated at power-up

Default

102 400 bytes

Use

unlockable_mem defines the minimum amount of memory which is guaranteed to be available for virtual memory and/or system overhead at any one time.

It limits the amount of memory which can be locked (lockable memory) to (the available memory indicated at powerup - unlockable_mem).

If unlockable_mem is greater than available memory, then none of the memory is lockable.

Lockable memory is used for:

- process images and overhead locked with *plock(2)*
- shared memory segments locked with the SHM_LOCK command of the *shmctl(2)* system call
- miscellaneous dynamic kernel data structures used by the shared memory system and some drivers.

Any call which needs to use lockable memory may fail if its value is too small. Note that lockable memory is a limit on the amount of memory which can be locked, but that this memory is available for virtual memory except when it is locked.

Cost

None.

Dependencies (interations with other system values)

`unlockable_mem` <= physical memory

using_array_size

Name

`using_array_size`—size of the cluster cnodes' using array

Range

0 to memory limited

Default

0	If the cluster code is not configured into the kernel.
`nproc`	If the cluster code is configured into the kernel. Unless you change other configurable parameters, this value will be 102.

The cluster code is configured into the kernel if the kernel description file you used contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

Use

`using_array_size` defines the size of the kernel's using array. The using array is an array of kernel structures that holds information related to outbound requests. It is the same size as the `nproc` configurable parameter.

Each outbound request consumes a single `using_array[]` entry that is held until the request is satisfied. Outbound requests are not discarded if all `using_array[]` slots are utilized; they are put to sleep until the required resource is available.

If both of the configurable parameters, `dskless_node` and `server_node` are equal to zero, the system is treated as a standalone, non-clustered kernel. In this situation, `using_array_size` is zero and `using_array[]` is not compiled into the kernel.

Cost

16 bytes × `using_array_size`

Dependencies (interactions with other system values)

The size of the using array is directly dependent upon the `nproc` configurable parameter.

Swap Space Parameter Interaction

If you change maxdsiz, maxssiz, maxtsiz, or shmmax, there are four other operating system parameters that you must also change. The four operating system parameters (dmmin, dmmax, dmtext, and dmshm) do not have a separate description page in this chapter, but Table A-1 explains the values these operating system parameters must have for given values of maxdsiz, maxssiz, maxtsiz, or shmmax. All of these swap space system parameters interact, and a wrong value might make your virtual memory system unworkable.

Note The system parameters dmmin, dmmax, dmtext, and dmshm should only be changed when it is necessary to change maxdsiz, maxssiz, maxtsiz, or shmmax. At all other times, these parameters should be left alone.

In the following example, assume you have a stack size equal to 90 Mbytes. Using the information in Table A-1, you would set dmmin to 16, dmmax to 2048, dmtext to 2048, and dmshm to 2048. Based on this information, you would need the following entries in your configuration description file (/etc/conf/dfile):

```
* tunable system parameters

maxdsiz     116*1024*1024
dmmin       16
dmmax       2048
dmtext      2048
dmshm       2048
```

Refer to the explanation of dfile in the kernel chapter of this manual.

Table A-1 shows the values that should be set in order to make the swap space parameters function correctly. The sizes shown for `maxdsiz`, `maxssiz`, `maxtsiz`, and `shmmax` are in Mbytes. The sizes shown for `dmmin`, `dmmax`, `dmtext`, and `dmshm` are in 4 Kbyte disk blocks.

Note that if you have a Model 310 or less than 20 Mbytes of swap space, you may want to reconfigure your kernel with the numbers shown with a superscript 1 (^{1}) for optimum use of your swap space.

All numbers shown with superscript 2 (^{2}) are the default settings.

Table A-1. Virtual Memory Parameter Table

If maxdsiz or maxssiz is:	If maxtsiz is:	If shmmax is:	dmmin should be:	dmmax should be:	dmtext should be:	dmshm should be:
<= 15 MB	<= 10 MB	<= 3 MB	16^1	256^1	256^1	256^1
<= 30 MB	<= 20 MB	<= 6 MB	16^2	512^2	512^2	512^2
<= 59 MB	<= 40 MB	<= 12 MB	16	1024	1024	1024
<= 116 MB	<= 80 MB	<= 24 MB	16	2048	2048	2048
<= 228 MB	<= 160 MB	<= 48 MB	16	4096	4096	4096
<= 448 MB	<= 320 MB	<= 96 MB	16	8192	8192	8192
<= 880 MB	<= 640 MB	<= 192 MB	16	16384	16384	16384
<= 1728 MB	<= 1280 MB	<= 384 MB	16	32768	32768	32768
<= 3392 MB	<= 2560 MB	<= 768 MB	16	65536	65536	65536
<= 6656 MB	<= 5120 MB	<= 1536 MB	16	131072	131072	131072

B

Federal Information Processing Standard

The U.S. Government has published an announcement for a Federal Information Processing Standard (FIPS 151-1, hereafter called FIPS). Based on on the POSIX standard IEEE Std 1003.1-1988, FIPS specifies the behavior of a system in areas where the POSIX standard permits divergent behavior. In three areas, HP-UX permits a wider range of behaviors than the FIPS. Based on the announcement of the FIPS published in the *Federal Register, Volume 54, no. 70, April 13, 1989,* the following sections explain the configuration of HP-UX so it conforms to the FIPS choices, pending its approval.

Restriction of the chown(1) Function.

The POSIX standard permits an implementation to allow users to give away their own files (as does System V) or to restrict this action to privileged users (as does 4.3BSD). HP-UX has the system administrator control this with the privileged group facility (see *getprivgrp*(2) in the *HP-UX Reference*). The FIPS makes the following statement, which requires the 4.3BSD behavior:

```
The implementation shall support the option
_POSIX_CHOWN_RESTRICTED.
```

You can configure HP-UX to behave this way by removing users from groups with the privilege *PRIV_CHOWN*. The default behavior of HP-UX grants this privilege to all users, so you need to revoke the privilege each time you start up the system. The simplest thing to do is to put the following line in `/etc/rc`.

```
setprivgrp -n CHOWN
```

Group ID of newly created files

The POSIX standard permits an implementation either to set the group ID of a newly created file either to the effective group ID of the creating process (as does System V) or to the group ID of the parent directory of the new file (as does 4.3BSD). HP-UX follows the 4.3BSD semantics if the set-group-ID bit of the parent directory is set, and follows System V semantics otherwise. The FIPS makes the following statement, which effectively requires the 4.3BSD behavior:

```
The implementation shall support the setting of the
group ID of a file (when it is created) to that of its
parent directory.
```

An HP-UX system can be configured to behave this way by setting the set-group-ID bit of all directories in the system. When HP-UX is installed, this bit is not set on directories, so this requires setting it once for all directories where the FIPS behavior is desired. Executing the following command as superuser will do this for the entire system; the command should be executed with no NFS mounts or RFA netunams in effect:

```
find / -type d -exec chmod g+s \;
```

In practice an administrator may choose to leave the setgid bit off for some directories which are not associated with any group, because the effective group ID of the creating process may be more meaningful for files in those directories. Examples of such directories include `/tmp` and `/usr/tmp`. This can be done by constructing a more complex `find` command, or by turning the set-group-ID bit off for those files after the `find` command, with a command such as:

```
chmod g-s /tmp /usr/tmp
```

This practice may not conform strictly to the FIPS.

Truncation of filenames during pathname resolution.

When a filename specified by a user is longer than the maximum supported by the file system, the POSIX standard permits an implementation either to truncate the name to the supported maximum (as does System V) or to give an `ENAMETOOLONG` error (as does 4.3BSD). HP-UX follows the 4.3BSD semantics for file systems that support long file names, and follows System V semantics for other file systems. The FIPS makes the following statement, which effectively requires the 4.3BSD behavior:

```
The implementation shall support the functionality
associated with the feature {_POSIX_NO_TRUNC}.
```

The system administrator can get this behavior by setting up all file systems to support long filenames (see the chapter named "Managing the File System" in the *HP-UX System Administration Tasks* manual).

C

HP-UX Partitions and Filesets

This appendix contains tables that show the partitions and filesets. This appendix does not list all the files in the filesets. Previous releases of this manual may also have shown the files in each fileset. Limitations in the amount of available space make it necessary to omit this information.

Compared to the 6.5 release, the partitions and filesets have been revised for the 7.0 release.. With this in mind, you should approach the 7.0 release as follows:

- If you install the 7.0 release of HP-UX and have no previous installed release, you need not consider the 6.5 release in any way. Just install the 7.0 release.
- If you install the 7.0 release and have a previous release running on your system, back up your applications and your personal directories. Then, install the 7.0 release. Restore your applications and personal directories later.
- If you are currently running the 6.5 release of HP-UX, and you plan to update to the 7.0 release, consider the following items with regard to the partitions and filesets:
  - Compare the list of partitions and filesets in the 6.5 release with those shown in this appendix for the 7.0 release. This will give you a feeling for the differences.
  - Read the README first documents that came with the 7.0 release before you install or update anything. These documents provide information you must accommodate.
  - During an install or update process, take time to examine the information in the files in `/etc/newconfig/Update_info` before you complete the installation or update.

The following sections have information about the partitions and filesets.

Relating Partitions, Filesets, and Files

The media for an installation or update contains partitions, filesets, and files as follows:

- A partition provides a logical means of organizing HP-UX into approximately 10 entities. A partition typically contains 5 to 20 filesets. While your installation or update media has partitions, your HP-UX system has no partitions.
- A fileset provides a physical means of organizing HP-UX files. The `/etc/filesets` directory contains a list of the filesets on your system. During an installation or update, you focus primarily on the filesets you want.
- A file is what you know a file to be on HP-UX. A fileset typically contains many files. Some of these files can have dependencies on files that lie in other filesets. In general, when you want a particular functionality, loading all the filesets in a partition will load all the files that have dependencies.

The Partitions

The 7.0 release of HP-UX has the following partitions:

Partition Name	Functionality
DATABASE	Allbase tools and Sample database.
DIAGS	CE Utilities for S300 and peripherals and Online diagnostic tools.
GRAPHICS	The Starbase family of graphics (with the graphics kernel system).
NETWORKING	Provides networking services used by HP-UX (ARPA, LAN, NFS, RFA, NFT, Virtual Terminal, X.25, IP and Programmatic Acess).
NLIO	Native Language I/O Core. See the filesets for languages.
NLS	Native Language Core System. See the filesets for languages.
OS_ADMIN	Utilities for system accounting and trusted systems. Driver writing examples, non-HP terminfo database, and the System Administration Manager (SAM).
OS_CMDS	Minimum command set, supplementary commands and programming environments with libraries, MS-DOS ™, remote printing, SCCS and RCS, communication utilities, and update tools.
OS_DOC	Online pages and utilities, and text formatting and processing utilities.
OS_KERNEL	2-user or multi-user license, kernel libraries and headers, core command set, cluster commands, and floating point accelerator tools.
PROG_LANG	Programs, tools, utilities, and compilers for programming languages: C, Pascal, FORTRAN, and Ada.
WINDOWS	Programs, tools, utilities, and libraries for the X11 and X10 windowing environments. Also, the Personal Applications Manager (PAM).

The Filesets in Partitions

The next several sections show the partitions and the filesets they contain.

DATABASE Partition

ALLBASE1.38	new
ALLBASE2.38	new
ALLBASE3.38	new
SAMPLEDB.38	still the same

DIAGS Partition

CE_UTIL.3-	was PCEUTIL
DIAGNOSTICS.3-	new

GRAPHICS Partition

AGP_DGL.38	still the same
DGL_SKEL.3-	still the same
FAFM_MIN.38	was AFA_FM
FAFM_SUPL.38	was PFA_FM
HPGKS.38	still the same
SBDL_BLD.38	was SBDL
SBDL_DEMO.38	was SBDL
SBDL_MAN.38	was SBDL
STAR_BLD.38	was PSTARBAS
STAR_DEMO.38	was DSTARBAS
STAR_MAN.38	was MSTARBAS
STAR_MIN.38	was ASTARBAS

NETWORKING Partition

ARPA.38	was LAN_AB and LAN_MAIL
ARPA_MAN.38	was LAN_MAN
LANLINK.38	was LAN_AB and LAN_NS
LAN_BLD.38	was LAN
LAN_MAN.38	new
NFS_INCL.38	still the same
NFS_MAN.38	still the same
NFS_RUN.38	was NFS_CMDS
NS_MAN.38	was LAN_MAN

NS_SERV.38	was LAN_NS
SRM.3	only on SE tape, still the same
VT3K.38	new
X25_COM.38	new
X25_IP.38	new
X25_PA.38	new
is obsolete	the NFS_MANC fileset
is obsolete	the LAN_MANC fileset

NLIO Partition

HLVTC_JPN.38	was HLVTCJPN
KFA_FM.38	still the same
NJWSERV.3-	still the same
NLIO_CHS.38	was NLIOCHS
NLIO_CHT.38	was NLIOCHT
NLIO_JPN.38	was HLIOJPN
NLIO_KOR.38	was NLIOKOR
NLIO_MIN.38	was NLIOCORE
NLWIDGET.38	still the same
NLX10_CHS.38	was NLX10CHS
NLX10_CHT.38	was NLX10CHT
NLX10_JPN.38	was NLX10JPN
NLX10_KOR.38	was NLX10KOR
NLX10_SUB.38	was NLX10SUB
NLX11_CHS.38	was NLX11CHS
NLX11_CHT.38	was NLX11CHT
NLX11_JPN.38	was NLX11JPN
NLX11_KOR.38	was NLX11KOR
NLX11_SUB.38	was NLX11SUB
ROMAN_JPN.38	new
SMPLX_CHT.38	new
SMPLX_JPN.38	new
SMPLX_KOR.38	new
STICK_CHS.38	new
STICK_CHT.38	new
STICK_JPN.38	new
STICK_KOR.38	new

NLS Partition

AMERICAN.38	was NAMERICA
ARABIC.38	was NARABIC
ARABICW.38	was NARABICW
CFRENCH.38	was NCFRENCH
CHINESES.38	was NCHINESS
CHINESET.38	was NCHINEST
DANISH.38	was NDANISH
DUTCH.38	was NDUTCH
ENGLISH.38	was NENGLISH
FINNISH.38	was NFINNISH
FRENCH.38	was NFRENCH
GERMAN.38	was NGERMAN
GREEK.38	was NGREEK
HEBREW.38	was NHEBREW
ICELANDIC.38	was NICELAND
ITALIAN.38	was NITALIAN
JAPANESE.38	was NJAPANESE
KATAKANA.38	new
KOREAN.38	was NKOREAN
NLS_CORE.38	was NCORE
NORWEGIAN.38	was NNORWEGI
N_COMPUTE.38	was NNCOMPUT
PORTUGUES.38	was NPORTUGU
SPANISH.38	was NSPANISH
SWEDISH.38	was NSWEDISH
TURKISH.38	was NTURKISH

OS_ADMIN Partition

ACCOUNTNG.38	was PACCT
AUDIT.38	was AAUDIT
DRIVERS.3-	new
INSTALL.3-	was INSTPROG
NONHPTERM.38	was ATRMINFO
PECC.3-	only on SE tape, still the same
SAM_MIN.38	new
TOOL.38	still the same

OS_CMDS Partition

CMDS_MIN.38	was ABCMD
DOS_UTILS.38	was PAPPLIC
KERN_SUPL.38	was PAPPLIC, PCORE, PFILTER, and PPERIPH
PRLP.38	still the same
PROG.MIN.38	was PPROG
PROG_SUPL.38	was PPROG
SRC_CNTL.38	was PSCCS
SYSCOM.38	was PSYSCOM
TOOL.38	still the same
USRCONTRB.38	only on SE tape

OS_DOC Partition

DOC_MIN.38	was AMANUAL
DOC_SUPL.38	was HPUX_MAN
TEXT_FMT.38	new
TEXT_READ.38	was ATEXT
TEXT_SUPL.38	was PTEXT

OS_KERNEL Partition

02_USER.3-	not known at release
DISKLESS.38	was PDSKLESS
DRIVERS.3-	was PDRIVERS
FPA.3-	was AFPA
KERN_BLD.38	was ACONFIG
MULT_USER.3-	not known at release
UX_CORE.38	was ACORE, CORE, SSIT, and KERNEL

PROG_LANG Partition

ADACMD.3-	new
ADADS.3-	new
ADAGKSL.3-	new
ADAGKSP.3-	new
ADASQLL.3-	new
ADASQLP.3-	new
ADASTARL.3-	new
ADASTARP.3-	new
C_MIN.38	was AC
C_SUPL.38	was PC
FORTRAN.38	still the same
FTN_LIBS.3-	was FTN_CORE
FTN_SUPL.38	was PFORTRAN
LANG_MIN.38	was ACAT
PASCAL.38	still the same
PAS_LIBS.3-	was PAS_CORE
PROG_MIN.38	was APROG

WINDOWS Partition

AGRM.38	still the same
PAM.38	was APAM
WIN_BLD.3-	was PWINDOW
WIN_DEMO.3-	was DWINDOW
WIN_MAN.3-	was MWINDOWS
WIN_RUN.3-	was AWINDOW
X10_BLD.38	was XWINDOWS
X10_DEMO.38	was XWINDOWS
X10_MAN.38	was XWINDOWS
X10_RUN.38	was XWINDOWS
X11BLDAUX.38	was PX11PENV
X11MANMIN.38	was AX11MAN
X11MANSUP.38	was PX11MAN
X11_BLD.38	was PX11PENV
X11_FONTA.38	new
X11_FONTB.38	new
X11_FONTC.38	new
X11_MIN.38	new
X11_RUN.38	was AX11RENV
X11_SERV.38	was AX11SERV

Index

D

G

H

I

K

L

M

N

O

V

W

X

HEWLETT
PACKARD

HP Part Number
98594-90061

Microfiche No. 98594-99061
Printed in U.S.A. E0989

98594-90661
For Internal Use Only