Contents

Contents

Contents

Contents

Module 8 — Disks and File Systems

Module 9 — File System Maintenance

Contents

Contents

Tables

Contents

Overview

Course Description

This 3 day class teaches students how to successfully perform the common system administration tasks necessary to maintain HP-UX systems that are dedicated to running technical or commercial applications. This is a lecture and lab course which emphasizes use of SAM as a system administration tool. Only those HP-UX concepts that are essential to make effective use of SAM, and other basic system administration commands, will be covered.

Student Profile and Prerequisites

This course is intended for system administrators/operators who need to perform the routine tasks necessary to maintain HP-UX systems that support application users. This administrator is not in charge of overall system configuration and overall network management. It is assumed that there is either another administrator (with more experience) or a third party support representative (from HP or other company) that will dictate the computing configuration and strategy.

Typically, students maintain HP-UX systems that are dedicated to running one or more technical (e.g. ME or EE CAD) or commercial (e.g. database, desktop publishing) applications. This class is aimed at the student who must know how to perform common system administration tasks but does not need to, nor want to, understand the details of how the HP-UX system performs the tasks.

The student is expected to be familiar with basic UNIX concepts such as directory structure, file permissions, groups, home directory, and processes.

The prerequisite for this course is "UNIX System Basics I" (51489S) or "Fundamentals of the UNIX System" (51434S).

Curriculum Path

```
                   |
    51489S         |
  UNIX System      |
    Basics         |
                   |        50722S                 H2572S
     or    ----->  |---->  HP-UX System       ---->UNIX System
                   |   Administration Basics      Basics II
    51434S         |
     UNIX          |
  Fundamentals     |
                   |
```

Overview

Notes to the Instructor

Introduction

KNOW THE STUDENT PROFILE! This is **NOT** a three day version of the 5-day system administration material! Nor is this class simply a class on SAM. It presents the essential and most common administration tasks using SAM as a delivery tool whenever possible.

This course is also more than just an operator course. Operators would not have anything to do with adding and configuring new disks into the system, for example. If you are dealing with an operator type audience topics that do not apply should be either left out or just skimmed.

Students with previous HP-UX system administration experience probably do not belong in this class. Students in charge of a network of systems or students in charge of any type of advanced system configurations should attend the 5-day class.

Note

The big distinction between the student intended for this class vs. the student in the 5-day administration class is that this student is responsible for *maintaining* a running, already-configured system. The student in the 5-day administration class is responsible for the *configuration, setup and overall management* of an HP-UX system.

Keep this in mind when delivering this material.

This is a practical, task oriented course which focuses on administration TASKS (how do I get the job done?). Information about what goes on in the HP-UX system when a task is executed is covered only when it is essential to understanding how to use SAM or if SAM cannot perform the required task. Many examples are given in the slides and student notes to help students see practical implementations of the material presented.

Whenever possible, this course will teach students the use of SAM to perform system administration tasks. The primary UNIX concepts covered are those required to understand how to use SAM effectively. The purpose of this course is to train individuals with minimal knowledge of the HP-UX system to successfully perform basic system administration tasks that are essential to maintain Series 300, 400, 700, and 800 HP-UX systems. In a few cases, more than the minimal amount of material is presented to appease students that may be more advanced than the "target" student.

The instructor kit is designed to assist you in preparing, setting up and presenting the *HP-UX System Administration Basics* class. Many advanced examples are given in the instructor notes if the students are more advanced than expected.

Notes to the Instructor

Instructor Profile and Prerequisites

The instructor for this class should have completed this class, the 5 day system administration course for both the Workstation and Series 800 (courses 51436 and 51482) or have equivalent knowledge. It is NOT required that the instructor have taught the 5-day courses. In fact, this course has many reference and preparation notes that will help an instructor not only prepare for this class but for both of the 5 day administration classes.

Agenda

The agenda is included *only* in the instructor's notes for this course. This is a suggested schedule for the course. You may adapt this schedule to meet your students' needs.

Daily Schedule

Day	Modules
Day 1 am	1, 2,
Day 1 pm	3, 4
Day 2 am	5, 6, 7
Day 2 pm	8, 9
Day 3 am	10, 11
Day 3 pm	Appendicies

The three appendicies are included as optional material. You may or may not choose to cover them.

If you are doing a "Required tasks only" class you will need to cover modules 1-6 and 10. Parts of the other modules not covered can be used to enhance what is presented. This would be a good flow for an "operator" class.

Classroom Setup

Hardware

The CRC (Central Registration Center) is supposed to ask students enrolling in this course what type of HP-UX system they have. This will aide you in setting up the classroom properly with the correct number of systems and seats on the systems.

This is the machine specific configuration information:

For Series 800:

- Two hard disks (minimum of 304 Mb disk space for the root disk any size will do for the second disk. Note that each lab team running the pseudo superuser environment requires about 3-4 Mb of disk space.)
- A reel tape or DDS/DAT drive for system backup.
- One terminal for system console
- One terminal for each student.
- A system printer.

There should be one workstation per two students. The recommended equipment for each workstation is:

For Workstations:

- Two hard disks (minimum of 300 MB disk space for "root" disk and any size for the second disk.).
- One monitor and keyboard (system console).
- Cartridge or DAT tape for backup and system recovery.
- A terminal for second student and to test adding a terminal.
- A printer for testing the lab on adding a printer to the system

Optional Hardware:

- LAN cable, Ts and terminators as needed to connect all systems.
- MUX cards for the workstations or a K2292 adapter for a Series 400.
- CD-ROM drives for installation/update media and/or LaserROM documentation databases.

Software

- 8.0 HP-UX installation/update tapes for Series 300/400/800s.
- 8.05 or 8.07 HP-UX media for the Series 700s.
- 8.02 HP-UX media for the 8X7 (if you have one).
- Authology simulator software for simulating system installation, update and shutdown in the Series 800 labs.(*)
- Pseudo Superuser environment software (for most other Series 800 labs).(*)

(*) These are taken from the lab tape for the 5-day HP51482 for the 800.

Student Account Configuration

If using the pseudo superuser environment you may choose to have a separate "normal" student account for each student but this is not required.

Notes to the Instructor

Preparation Tasks

See the information noted in the Instructor Profile section.

The key to teaching system administration is to DO IT. You cannot effectively learn system administration without experiencing it. Before you teach this class make sure you have taken the 5 day administration courses. You must also work with *this* course material. You must read it thoroughly, try the examples, and *do the labs*. This course is very different from the other system administration courses. Do not be caught off guard by not working through this material before teaching it.

Materials List

Instructor Materials

- Instructor kit (P/N 50722-60008)
- Overhead slide kit (P/N 50722-60009)
- HP Support Services Instructor Module (P/N 5960-2978).
- HP Support Services Instructor Overhead Slide Kit (P/N 5960-2979).

Student Materials

- Student kit (P/N 50722-60007)
- HP Support Services Student Module (P/N 5960-2977).

Classroom Materials

Library List

The following reference manuals should be available in the classroom:

- HP-UX System Administration Tasks Manual.
- HP 9000 HP-UX Reference
- Installing and Updating HP-UX
- Installing Peripherals

Current part numbers (as of printing) are given in module 1.

Having a few systems configured to run the LaserROM software is also recommended and very useful in the classroom.

Presentation Materials

- Whiteboard
- Overhead projector
- Extra transparencies and felt tipped pens

Module 1 — Introduction to System Administration

Objectives

Upon completion of this module, you will be able to do the following:

- Identify responsibilities of a system administrator.
- Describe the topics covered in this course.
- Identify the HP-UX documentation and tools targeted towards helping system administrators.

About This Module

In this module we describe the responsibilities of an HP-UX system administrator in an application environment. In addition we present an overview of the system administration tasks and tools covered in this course.

Note All references to Series 400 systems also apply to Series 300 systems and all references to Series 800 systems also apply to Series 600 systems except where noted otherwise.

Module 1 — Introduction to System Administration

Overview of Module 1

We want to identify why the students are here and the reference manuals they will need to accomplish their jobs.

1-1. SLIDE: Welcome to HP-UX System Administration

Welcome to HP-UX System Administration

- Introductions
- What do you want to get from this course?
- Logistics
- Prerequisites
- Organization of the Student Guide

50722 1-1 1

Student Notes

This course is designed to teach you enough information to become an effective administrator of an HP-UX system that runs a turn-key application, and to teach you how to learn more on your own.

The student guide is designed for use both in this class and later as a reference tool. There is a Table of Contents that provides an outline of the course topics. Note that most topic pages have a **SLIDE** and additional information in the section entitled **Student Notes**. If you have time, read through the Student Notes before each lecture. This will help you get the most out of your classroom time. After the lecture, use the Student Notes to reinforce and review what you have learned.

1-1. SLIDE: Welcome to HP-UX System Administration

Instructor Notes

Purpose

Before digging into the course material, take some time to get acquainted with your students and share logistical information with them. We all have our own Monday morning class-room intro act but the following list is offered as a minimal checklist of topics to be covered.

Key Points

- Introduce yourself to the class and discuss your background and qualifications.
- Have the class members introduce themselves and describe what they want to get out of this course.
- Explain how you will run the course, for example, class time, breaks, the degree of formality you would like, etc.
- Inform the students of logistical details such as:
  - phones
  - photocopying machines
  - rest rooms
  - coffee and pop/soda machines
  - where to take breaks and lunch
  - the local guidelines for smoking
- Remind students that the prerequisite for this course is the two day course "HP-UX Basics for Application Users" (51489B), although the five day courses "Fundamentals of HP-UX" (51434B) is also an acceptable prerequisites.
- Students should be familiar with basic UNIX concepts such as directory structure, file permissions, groups, home directory, processes, and editing files. (You might want to review these concepts if you feel it is needed.)
- Explain how the student guide is organized.

Teaching Tips

It is important to get somewhat acquainted with your students before you begin teaching. By having the students introduce themselves, it gives you a chance get a feel for the group, and it also gives the students a chance to settle in, find out who else is in the class, and have their basic questions answered.

Ask your students to introduce themselves by providing:

- name
- company
- job responsibilities
- hardware at their site, for example which SPUs they have (this is important so you can spend class time where the students need it most and you will not be covering irrelevant material).
- applications installed on the hardware
- UNIX background and experience
- expectations of this course and what they want to get out of this course (you can write their expectations on a flip chart and refer to it later as topics are covered).

It is your job to make this course and this week a positive experience for the students. Many times this means sharing information other than course material. Tell your students about great places to visit during their off hours or your favorite restaurant in town. Do your best to help make their stay in your city enjoyable.

Point out that students who complete the curriculum path for HP-UX application users may pursue the curriculum path created for individuals who wish to acquire a deeper understanding of HP-UX operating system concepts and to learn how to perform a wider range of system administration tasks.

For students from a commercial environment, the path for system operators is *UNIX Basics I* (51489B) and *HP-UX System Administration Basics* (50722C). It is common in the commercial environment for experienced system operators to move into the role of system manager after a period of time. The growth path for system operator to system manager is to attend *HP-UX Basics II*, followed by the appropriate 5 day system administration course.

Transition

Let's talk about the role of the system administrator.

1-2. SLIDE: The Role of the HP-UX System Administrator

The Role of the HP-UX System Administrator

Responsibilities to the:

- users
- software
- hardware

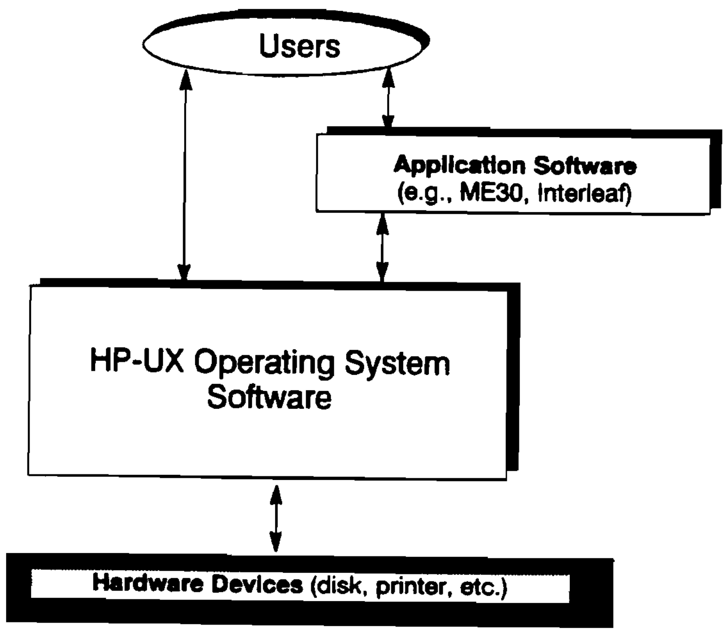

50722 1-2 2

Student Notes

The responsibilities of the system administrator may be divided into several categories:

- Understanding the needs of the user community.
- Managing the system's hardware.
- Managing the system's software (both application software and HP-UX operating system software).

Since many of the tasks associated with these responsibilities require access to commands that should not be available to everyone, the system administrator needs special access to the system. This access is called "superuser" or "root" access.

Responsibilities to the User

The system administrator is responsible for creating, removing and updating user accounts as well as maintaining the membership of groups.

Users frequently approach administrators with questions such as, "How do I . . . ?" or ask for help because "My terminal is broken" or "I forgot my password." This course will help prepare you for these types of questions.

The system administrator should understand the intended use of her or his computer system and be aware of the number of users on the system, the characteristics of each user, the system resources and peripherals required by each user, and the data/programs that must be shared by various user groups.

Responsibilities to the Hardware

The system administrator of an HP-UX system is responsible for verifying that the HP 9000 system hardware has been installed correctly and thoroughly tested. This must be done before an HP-UX operating system can be loaded. In some cases (especially on the Series 800 and 600), the administrator may not be the person who actually installs the hardware. Often a Hewlett-Packard Customer Engineer will perform the installation.

Once the system is operational, the administrator must monitor the performance of the various hardware components. If a hardware failure occurs, the administrator should attempt to isolate the problem to a component level. Depending upon the service agreements in place, the administrator may schedule a customer engineer to make necessary repairs.

The system administrator should be prepared to add peripheral devices such as printers and terminals to an HP-UX system.

Responsibilities to the Software

The system administrator is sometimes called upon to install and periodically update the HP-UX operating system. (In some cases, depending on which HP-UX system is purchased, the operating system may be installed by Hewlett-Packard personnel or is pre-installed at the factory). HP-UX systems that do not have the operating system pre-installed are shipped with the necessary media to boot and install the operating system. Several different media options are available dependent upon which model of SPU you have.

9-track reel-to-reel tape
16-track cartridge tape
4-mm DDS/DAT tape
CD-ROM disc

Once the HP-UX system has been installed, it may be necessary for the system administrator to install additional application software or to reconfigure the system.

We'll see later that HP-UX organizes its data on a disk using a file system. Since the file system is where all of the system and user data is stored, it is quite important to ensure that the integrity of the file system is maintained. This includes tasks such as:

- periodically backing up files to secondary storage (like tape)
- possibly restoring lost or corrupted files from the backup
- periodically checking for and correcting any file system corruption
- monitoring free space in the file system

Most HP-UX systems include output devices such as printers and plotters. The system administrator must manage the software which sends output to the printing devices.

Periodically Hewlett-Packard will release an update to HP-UX. An update may modify existing system features and/or add new capabilities. The system administrator is responsible for installing each software update so the HP-UX system available to the user community contains the latest version of the software.

Note It is also your responsibility to coordinate releases of HP-UX with the releases of your application software. Always check with the application supplier for revision compatibility before updating HP-UX to a new release.

1-2. SLIDE: The Role of the HP-UX System Administrator

Instructor Notes

Purpose

To provide an overview of system administrator responsibilities. (In the following sections of this module we derive the list of course topics from these more general system administration responsibilities.)

Key Points

- System administration encompasses all tasks related to the set-up and maintenance of an HP-UX system.
- A system administrator is a superuser who has capabilities not available to regular users.
- It is up to the system administrator to monitor the performance of the hardware and peripheral devices. When a user has a problem with a terminal or a printer, or is running out of disk space, he/she will generally come to the system administrator for help. Be prepared!
- If a problem occurs, the system administrator should take action. If the administrator is not able to fix the problem, he/she should schedule a Hewlett-Packard Customer Engineer to make the necessary repairs.
- The system administrator must add new users to the system.
- The system administrator must tailor the system to fit the needs of the user community.
- Users will expect the system administrator to answer questions and solve their system problems.

Teaching Tips

Note

Emphasize that this course is geared toward individuals who administer one or a few HP-UX systems that are typically dedicated to running one or more technical (e.g., CAD) or commercial (e.g., database, desktop publishing) applications. Also remind students that this course will focus on teaching them how to do things and not delve deeply into how the HP-UX operating system works. This course is not intended for administrators of, say, 50 HP-UX systems networked together. Students who maintain complicated systems and/or want to know more about how the HP-UX system works should take the 5 day system administration course and possibly go onto the network administration course.

Ask your students to name some system administrator responsibilities and identify the category (hardware, software, user) that the tasks fit in. List the examples generated by the students on the board (or flip chart) under each category. Then fill in any additional tasks that the students may have left. This exercise will get students to start thinking about the content of the course and will lead into the next section of this module.

Transition

Let's look at the topics (that is, the system administration tasks) covered in this course.

1-3. SLIDE: Course Topics

Course Topics

- Module 1—Introduction to System Administration
- Module 2—System Hardware and Peripherals
- Module 3—System Startup and Shutdown
- Module 4—The System Administration Utility
- Module 5—User Accounts
- Module 6—Printers and Plotters
- Module 7—Terminals and Modems
- Module 8—Disks and File Systems
- Module 9—File Systems Maintenance
- Module 10—Backing Up File Systems
- Module 11—Other Administration Tasks
- Appendix A—Installing HP-UX
- Appendix B—Updating HP-UX
- Appendix C—Setting up a Local Area Network

50722 1-3 3

Student Notes

This course is designed to teach you the basic concepts and skills needed to administer an HP-UX system. By no means do we cover everything there is to know about system administration, but what you learn in this class will get you going and give you a solid base upon which to build more knowledge.

The sequence in which the course topics are introduced does not reflect the manner in which you would set up a new HP-UX system (beginning with hardware set up, proceeding to HP-UX installation, etc). Instead we try to introduce concepts so that the later modules build on or utilize the concepts introduced in the earlier modules.

1-3. SLIDE: Course Topics

Instructor Notes

Purpose

This slide will give the students an idea of what topics will be covered in the course and how the course will flow.

Key Points

- Point out the modules and appendices that correspond to the topics shown in the slide. You may want to tell students which modules will be covered on which days of this course and give a very high level explanation of what each module will cover.

Teaching Tips

- You may want to redraw the slide on a flip chart and refer to it as a road map to the course topics.

1-4. SLIDE: System Configuration Files

System Configuration Files

When the HP-UX Kernel has to answer the basic question of:

How?
Where?
When?

It gets its answers from the configuration files:

```
      /
      |
                    -------|-------
     etc
      |
           ---------------------------------
      |        |         |          |
          passwd...  groups...  hosts...  inittab
```

And many many more!

50722 1-4 4

Student Notes

The HP-UX Operating System (referred to herein as the "O/S" or the "kernel") is very flexible and configurable. These are very admirable traits but require the O/S to be configurable by the end user (you!). This configurability is accomplished by storing configuration information in simple ASCII files which are commonly referred to as databases or tables.

Administration can be performed by carefully editing the information in these files. Long-time UNIX gurus often use this approach. Over the years many new utilities have been created to make administration tasks easier. In this course we will try to show you the higher level utilities available for your use, and we will occasionally have to take a look at some of the "data bases" underlying certain tasks.

If you need to know how certain configuration files are structured, refer to the "File Formats" section of the *HP-UX Reference Manual* (section 4).

These files have their permissions set for read-only access. This reduces the chances of accidently changing or overwriting the file's data. This means that only the superuser (root) can make changes to these files. If you must make changes to a configuration file manually, you must log in as root and edit the file with a text editor (such as `vi`). When you try to save your changes, you will be warned that the file is read-only. At this point you will need to override the read-only protections from within the editor. In the `vi` editor this is accomplished using the command

```
:wq!
```

(read colon w q bang). This will save your changes and exit the editing session.

1-4. SLIDE: System Configuration Files

Instructor Notes

Purpose

We would like to initiate the student to the concept of configuration files.

Teaching Tips

- Don't dwell on the syntax of any particular file at this point. Just point out some of the "databases" found on a typical system.
- Try not to spend too much time here as several of the databases will be covered in later sections.

Transition

Which HP-UX commands are needed in order to perform the administration tasks we have identified?

Handout: `vi` Cheat Sheet

Since we will probably be using `vi` often in this course, students may need a refresher on its basic functions. The next page can be photocopied for the students.

vi Editor Quick Reference

How to use this chart:

1. Type the commands as they are shown in the second column below.
2. Follow the colon (:) and search (/) commands with Return as indicated.

To Do This ...	Type This in Command Mode
Getting Started with vi:	
Read, create, or edit a file (*enter from shell*)	`$:vi` *filename*
Make sure you are in command mode	ESC
Show whether you are in command or insert mode	`:set showmode` Return
Find out what file you are in	CTRL + g
Moving Within the File:	
Move the cursor left or right	Use `h` or `l` key
Move the cursor up or down	Use `k` or `j` key
Cursor to end of line	`$`
Cursor to beginning of line	`0` (zero)
Cursor to end of file	`G`
Cursor to beginning of file	`1G` (one G)
Editing Text — Inserting:	
Insert text at the cursor position	`i`
Insert (append) text *after* the cursor position	`a`
Backspace to overwrite previous character (*in Insert Mode*)	CTRL + h or Backspace
Editing Text — Deleting:	
Delete characters at cursor	`x`
Delete word (cursor at beginning)	`dw`
Delete entire line at cursor position	`dd`
Delete from cursor to end of current line	`D`
Delete from cursor to end of file	`dG`
Moving Text:	
Join lines of text	`J`
Copy a line into a buffer	`yy`
Put copied or deleted text line after cursor line	`p`
Read in another file after cursor line	`:r` *filename* Return

vi Editor Quick Reference (Continued)

To Do This . . .	Type This in Command Mode
Searching and Marking Text:	
Search for *words*	/*words* (Return)
Repeat the previous search for *words*	n (forward) or N (backward)
Place invisible mark at current cursor position	m (*a-z*)
Move cursor to invisible marker	: (*'a-'z*) (Return)
Saving and Printing Files:	
Save file in same file name	w (Return)
Save file to a new file name	w *new_file_name* (Return)
Exit from file without saving changes you've made	q!
Overwrite another existing file with this file	w! *another_file_name* (Return)
Save and exit from file	wq (Return)
Print a file from within that same file	:!lp % (Return)
Print another file from within a file	:!lp *file_name* (Return)
Repairing Mistakes:	
"Undo" the previous action	u
Restore a line to its previous state	U
Restore ("put") last delete	p
Restore current file, disregarding changes since saving	:e! (Return)
Compare changes "before and after" an action	u *(repeat to "toggle")*
Recover a file after a system interruption (from shell)	$:vi -r *filename*
In General:	
Execute shell command *cmd* from within file	:set wm=*n* (Return)
Set line length (wrapmargin) *n* lines from right	:set (Return)
Display current options	:set all (Return)
Display all options	:ab *str string* (Return)
Abbreviate *string* with *str*	:ab (Return)
Display all abbreviations	:una *str* (Return)
Unabbreviate *str*	

For more information, see *The Ultimate Guide to the vi and ex Text Editors*.

1-5. SLIDE: Overview of System Administration Commands

Overview of System Administration Commands

Major HP-UX commands used in this course:

- `sam` for most administration tasks
- `shutdown` and `reboot` to stop (restart) the system
- `wall`, `news` and `elm` for communicating with users
- `passwd` for changing user passwords
- `lpadmin`, `disable`, `enable` and `cancel` for print spool control
- `fsck`, `mount`, `umount` and `newfs` to build and maintain disks
- `fbackup`, `frecover` and `tar` for file system backups
- `cron` and `crontab` to schedule jobs for automatic execution
- `update` for loading and updating system software

50722 1-5 5

Student Notes

This slide lists a number of the system administration utilities that you will learn to use during this course. Some, like `sam` you will use quite often. Others, will be used much less frequently (like using `update` to load new system software).

1-5. SLIDE: Overview of System Administration Commands

Instructor Notes

Purpose

To summarize the major HP-UX commands that students will learn in this course.

Key Points

- Briefly review the HP-UX commands that will be covered in the modules and appendices as a lead into the manuals that system administrators use.
- As the slide indicates, SAM can perform a variety of system administration tasks.

Transition

Besides the SAM program you will frequently use other commands to help you administer your system. The HP-UX documentation set describes the usage of all of these commands. Lets look at these manuals.

1-6. SLIDE: HP-UX Administration Manuals

HP-UX Administration Manuals

Required System Administration Manuals:

- *Finding HP-UX Information*
- *Master Index*
- *HP-UX Reference*
- *System Administration Tasks*
- *Installing Peripherals*
- *Installing and Updating HP-UX*

Other good manuals to have:

- *How HP-UX Works: Concepts for the System Administrator*
- *Error Message Catalog*
- *Managing Clusters of HP 9000 Computers*
- *Solving HP-UX Problems*
- *HP-UX System Security*

50722 1-6 6

Student Notes

Many times when you use your system, you will need to look something up in the documentation. There are many HP-UX manuals. How do you know which manual to check? There are several HP-UX manuals that are targeted specifically towards administrators. You should be familiar with these manuals and their content as they can be of tremendous help to you. These manuals are listed below.

Note that some of the HP-UX manuals are specific for one platform (300/400, 700, or 600/800), but many of the manuals are identical for the entire HP 9000 family. The entire manual set is available on CD-ROM as well and is accessed with a tool called LaserROM/UX.

Finding HP-UX Information

Three versions of this manual exist, one for each platform. This manual is designed to help you locate information that will help you work with the HP-UX operating system. There are many HP-UX manuals available, and you may find it difficult to decide which one to use. This manual lists all of the HP-UX documentation available with the part number, intended audience, and tasks and concepts covered. If you look through this manual, you should be able to figure out what other manuals you need.

Master Index

Three versions of this manual exist, one for each platform. This index will refer you to the appropriate HP-UX manual and allows you to look things up by *action*, *object*, or *feature*. It also provides a table of the full manual set by title and part number.

HP-UX Reference

This manual set (three volumes) is intended for all HP 9000 HP-UX systems. It is intended as reference material and is most useful to experienced users. It is not designed to serve as a learning tool for beginners. The manual set is comprised primarily of HP-UX command descriptions and supporting information. Remember that the content of this manual set is also available on-line by using the `man(1)` command.

The HP-UX Reference Manual is divided into eight sections:

- Section 1: User Commands
- Section 1M: System Administration Commands
- Section 2: System Calls
- Section 3: Subroutines
- Section 4: File Formats
- Section 5: Miscellaneous Facilities
- Section 7: Device Files
- Section 9: Glossary

This course will draw heavily upon section 1M, System Administration Commands. Section 1M contains information on those commands that are used primarily by a system administrator. Section 4, File Formats, is also invaluable to the system administrator as it contains information on most of the configuration files that you will be responsible for maintaining.

System Administration Tasks Manual

Three versions of this manual exist, one for each platform, but the layout and much of the content is identical in all versions. These manuals cover most system administration tasks in detail. Topics covered include constructing and customizing an HP-UX system, updating HP-UX, starting and stopping HP-UX, setting up and administering backups and the LP spooler, managing run-levels, users, groups, file systems, and HP-UX clusters.

Installing and Updating HP-UX

Three versions of this manual exist, one for each platform. The manual covers how to install, update and modify the HP-UX operating system. Updating HP-UX from another system across the network and Series 700 network installation are also covered in detail.

Installing Peripherals

Three versions of this manual exist, one for each platform. These manuals provide step-by-step instructions on how to configure peripherals, such as terminals, printers, plotters, as well as disk and tape drives.

Error Message Catalog

This manual is common to the HP 9000 HP-UX family. This guide helps to identify where a specific error message comes from and some possible corrective actions. Most notable are the messages from the kernel and from the file system maintenance tool, fsck.

How HP-UX Works: Concepts for the System Administrator

This manual supplements the *System Administration Tasks* manual to describe the essential concepts of system administration. Whereas the *System Administration Tasks* manual describes how to administer your HP-UX system, this manual discusses the underlying principles.

It explains HP-UX system concepts useful to programmers, users and system administrators. Concepts covered include system startup, system shutdown, login, process control, run-levels, memory management, file system, device files, system configuration, using peripherals, networking, and system accounting.

E/ISA Configuration Guide for HP-UX

This guide explains how to configure EISA and ISA (E/ISA) I/O boards on Series 400 and 700 workstations containing EISA slots. It describes the interactive program eisa_config(1M), which has been provided to help you with this configuration.

HP-UX System Security

This manual is common to the HP 9000 HP-UX family and contains procedures and guidelines essential for maintaining a secure HP-UX system. The manual provides detailed information on implementing HP's full range of security features, such as auditing and refined control mechanisms for file access.

Managing Clusters of HP 9000 Computers

Three versions of this manual exist, one for each platform. This manual covers the steps necessary to create a clustered environment and explains how to successfully administer an HP-UX cluster. Other topics cover how the sharing of the HP-UX file system between client nodes and server nodes is done.

Solving HP-UX Problems

This manual is common to the HP 9000 HP-UX family This is a relatively new manual and helps the system administrator identify and solve some of the more common administration headaches.

System Support Log (for Series 600/800)

This binder contains all of the forms necessary to record the complete history of an HP-UX computer system. Here you can collect logging reports about preventive maintenance activities, records of hardware and software revision changes, and I/O configuration maps.

Table 1-1. System Administration Manuals for HP-UX Release 8.X

Manual Title	Part Number Series 300/400	Part Number Series 700	Part Number Series 600/800
Finding HP-UX Information	B1862-90001	B2355-90007	B2437-90001
Master Index	B1862-90003	B2355-90008	B2437-90004
HP-UX Reference	B1864-90000	B2355-90004	B1864-90000
System Administration Tasks	B1862-90008	B2355-90003	B2437- 90006
Installing and Updating HP-UX	B1862-90002	B2355-90000	B2437- 90003
Installing and Updating HP-UX 8.07		B2355-90014	
Installing Peripherals	B1862-90007	B2355-90006	B2437-90005
Error Message Catalog	B1862-90004	B1862-90004	B1862-90004
How HP-UX Works: Concepts for the System Administrator	B1862-90005	B2355-90005	B1862-90005
E/ISA Configuration Guide for HP-UX	B2370-90000	B2355-90012	
HP-UX System Security	B1862-90009	B1862-90009	B1862-90009
Managing Clusters of HP 9000 Computers	B2355-90009	B2355-90009	B2355-90009
Solving HP-UX Problems	B1862-90010	B1862-90010	B1862-90010

LaserROM allows users to quickly search through the entire set of HP-UX manuals for specific topics. This is an extremely useful tool that will be more widely used in the near future. To use LaserROM you need a CD-ROM drive (available from HP) and a LaserROM subscription for HP-UX. This subscription updates your entire manual set (which is on one CD-ROM disk!) every month.

1-6. SLIDE: HP-UX Administration Manuals — Instructor Notes

Purpose

It's important that the students learn which manuals are useful to their job as system administrator and where to find the appropriate information in these manuals.

Key Points

- The administrator cannot do the necessary job in a vacuum. There must be manuals and reference materials available.
- Amongst these materials the *HP-UX Reference* and the *HP-UX System Administration Tasks Manual* are the principal reference manuals needed by an administrator.
- Explain the value of the *System Support Log*. Proper updating of the records in this binder not only eases system reconfiguration or trouble-shooting, but is also essential if two or more people are responsible for the system administration. The *System Support Log* is the best place to share information related to their customized system.
- Recommend other reading and reference materials.
- We dare anyone to figure out our manual part numbering scheme! They constantly change and go out of date. One of the best places you can find up-to-date manual part numbers is in the monthly LaserROM table of contents.

Teaching Tips

- You should constantly reference the appropriate manuals during this course. This will allow the students to become familiar with the manual names and sections. Have the students look up particular concepts frequently to get them used to using the manuals.
- Go over the **Reference** notes when presented in the various topics.
- The students should have the *HP-UX Reference* and the *System Administration Tasks Manual* at their disposal. Go through these manuals with them and point out the sections that are relevant to them as administrators. Also, you should encourage the students to use the `man` command.
- If you have your documentation on LaserROM, take the time to show the students how LaserROM works and how they can access information. LaserROM allows users to quickly search through the entire set of HP-UX manuals for specific topics. Even if you don't have LaserROM, use this opportunity to tell students that it is a useful tool that will be used more widely in the near future. One physical advantage is that the entire set of HP-UX manuals can fit on one CD-ROM. Also, HP distributes LaserROM updates so students never need to worry about inserting update pages in their physical manuals.

Transition

Now that we have a feel for the task at hand, let's talk about system administration.

Module 2 — System Hardware & Peripherals

Objectives

Upon completion of this module, you will be able to explain or describe the following aspects of your HP 9000 Computer Systems hardware:

- Organization of the I/O system
- Built-in interfaces, types of connectors, and their locations
- Connection, addressing and power up of peripheral devices
- The naming convention of device files for the system
- Hardware addressing scheme of the system

About This Module

This module examines the hardware and the I/O (Input/Output) system of Series 300, 400, 700 and 800 computers. We describe the basic computer system architecture and provide examples of the interface cards and peripherals found on the various models.

We also describe one of the most important concepts dealing with peripherals, device files. This is how we will reference and use all peripherals in our system. This will be key to all the work we do with peripherals in the rest of the course.

This module does not intend to make you a peripheral configuration guru. It presents the terms, names and conventions used when dealing with devices on an HP-UX system. Most importantly, it gives you manual references and the basic knowledge you will need to use the manuals effectively to configure any type of peripheral on your system.

Overview of Module 2

Teaching Tips

- Our goal in this module is to get enough of a hardware overview so students can intelligently identify and use basic peripherals.
- Be somewhat aware of your student's hardware configuration before you start this section. You should probably cover all of the material but only spend time on what applies to the students. For example, if the entire class uses SCSI devices, you will not have to spend much time (if any) on HP-IB. Speak directly to the class's needs.
- Don't try to cover all of the variations of interface cards, peripherals, or SPU types. The slides and notes bring out the major points regarding the different systems where necessary. Trying to discuss the many variations just in the Series 800 family would take far too much time and just confuse the students.
- There are sufficient instructor notes to give more detail if you have an advanced audience. Do not feel as if you should present all of the student AND instructor notes. Most audiences will be blown away. Adjust the lecture content to the audience's needs.
- When discussing the various SPU types you should have the supplemental color hardware slide set to help describe the SPU layouts. If you did not get these slides with either of the 5 day system administration classes they are available through SRDO [P/N 5960-6653].
- The **Definition** and **Reference** notes should be pointed out. Definitions are given for terms as they are used (like DIO and EISA). References are given for students to learn more or find more specific information about their configurations. This chapter was designed to satisfy the needs of the system operator and administrators of turn-key application systems. The references will point the student in the right direction if they are more advanced.
- The details of major and minor numbers were purposely left out of this chapter. It is questionable if it is really needed for the target audience and there is no doubt that it is an extremely confusing topic to cover to an audience of so many different backgrounds. If you feel it is essential, feel free to do some "instructor value-added" work.
- Hardware seems to get out of date faster than courses can be updated. Make sure you check the most current configuration and pricing guides (available from literature distribution centers) for latest SPU releases and configuration options. You may want to photocopy newer material to supplement the course material.

2-1. SLIDE: I/O Architecture and the SPU

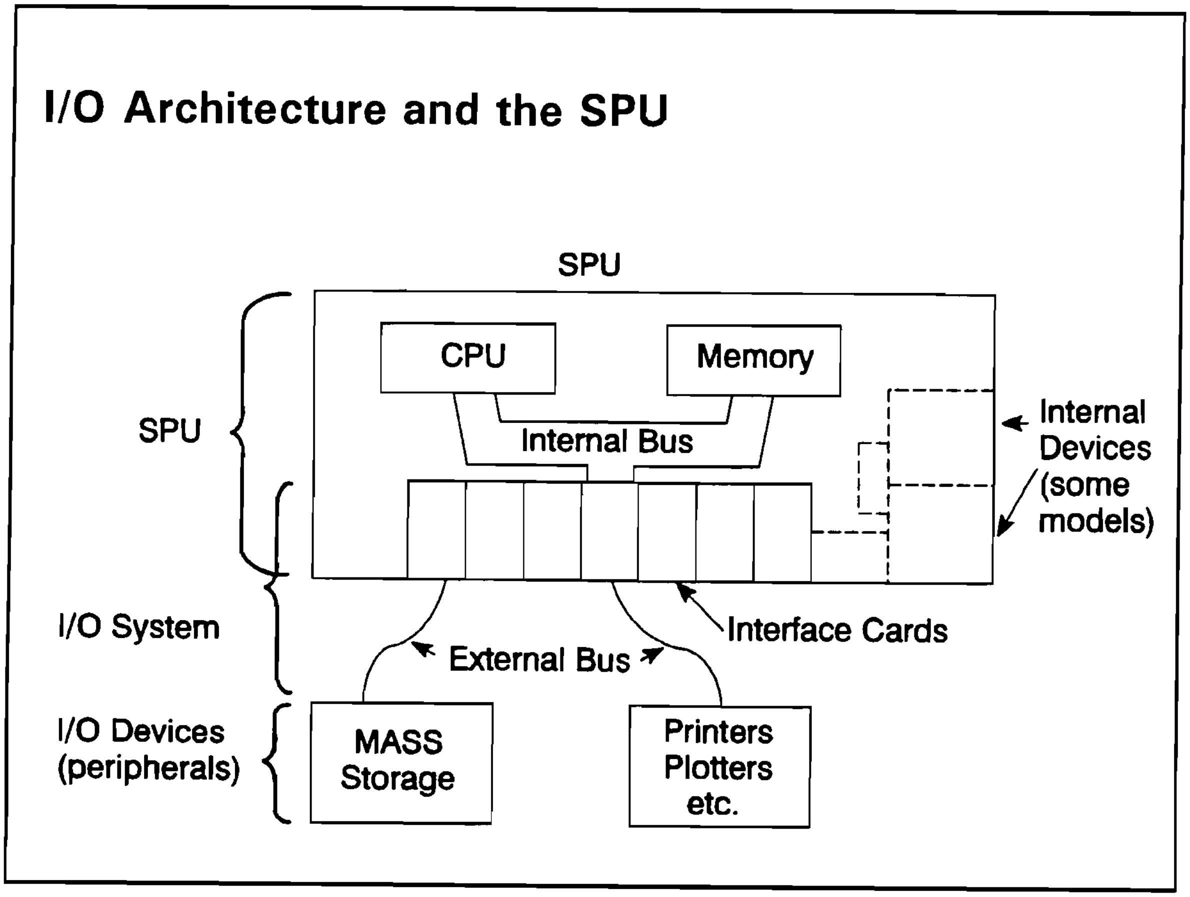

50722 2-1 7

Student Notes

The various members of the HP 9000 Computer Family utilize several different hardware architectures and I/O subsystems. The point to remember is that all the systems are designed to accomplish the same goal; transfer data between the central processing unit, the physical system memory, the virtual system memory (swap space), and the peripheral devices.

These transfers are accomplished via one or more internal and external buses. Buses are similar to streets that move traffic through a city, a particular location on a bus has an associated address just as your home has a street address. The hardware addressing varies in complexity and length as we move from the simple architectures to the larger, more advanced architectures.

Terminology

The following describes some of the components that are involved with system I/O:

SPU	System Processor Unit. The cabinet containing the CPU, system memory, and the I/O system. Many models will also have internal peripheral devices such as disk and tape drives. The SPU will also have one or more internal interface cards providing the mechanism to connect many types of external devices.
CPU	Central Processing Unit. In this slide, the term refers to the instruction-processing module inside the computer, not to the computer itself. The CPU processes data supplied to it by the I/O system.
Memory	Physical Random Access Memory (RAM) located in the SPU and available for use by the CPU and I/O System. This is where all data is operated on.
Internal Bus	The electronic path that connects the various areas of the SPU and allows data to flow throughout. Some models have several different internal buses joined by **bus converters** to facilitate efficient data flow between areas that operate at different speeds. (Similar to road systems where you find expressways, local main streets, and residential streets with the appropriate interchanges.)
I/O System	The physical hardware in the SPU that allows both builtin and add on interface cards to be connected.
Interface Card	Accessory card that is either built into the SPU or plugs into slots in the I/O system (typically on the back of the SPU case). This card will provide connections for and access to peripheral devices.
External Bus	A mechanism (like a cable) that can connect many like devices to one interface card. This is how devices like disks and tape drives are attached to the system.

2-1. SLIDE: I/O Architecture and the SPU

Instructor Notes

Purpose

To describe those components of the SPU that are involved with system I/O.

Transition

Let's categorize the basic parts of the computer system

2-2. SLIDE: Hardware Components of a Computer System

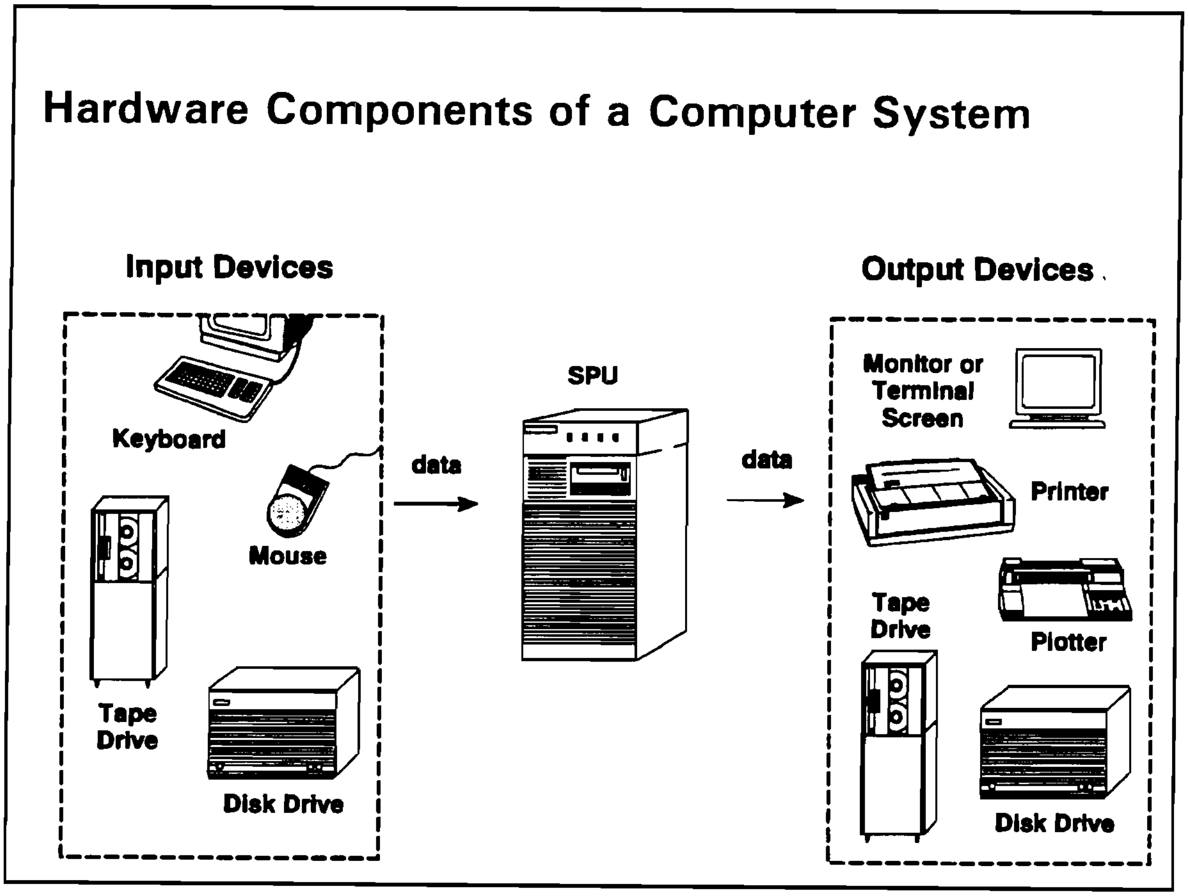

50722 2-2 8

Student Notes

An HP-UX computer system typically includes the following hardware components:

- The **SPU** itself.
- **Input devices** that accept data for processing or storage, including:
  - A **disk drive** that the SPU reads data from, transferring the data into memory for processing. (Many models have disks that are in the SPU.)
  - A **tape drive** which reads data from cartridge, DDS/DAT or reel tape and copies the data to memory and/or the system disk where it is stored for later use. (On some models the tape drive may be internal to the SPU.)
  - A **keyboard** or **mouse**, connected to a monitor or terminal, that is used to issue commands to the SPU or to input data that will be stored on disk.

- **Output devices** such as:
  - A **terminal** or **monitor** that displays messages generated by the SPU or data stored on the system disk.
  - A **printer** or **plotter** that creates a hardcopy of data generated by the SPU or stored on the system disk.
  - A **tape drive** which takes data stored on disk or in memory and copies it to a cartridge, DDS/DAT or reel tape.
  - A **disk drive** that the SPU writes data to for storage.

Note that I/O devices are also called peripheral devices.

2-2. SLIDE: Hardware Components of a Computer System

Instructor Notes

Purpose

To introduce some of the types of devices that the students will be dealing with on their systems.

Key Points

- The slide is meant to be generic and not represent any particular type of device. Students should not be worried if their SPU or tape drive doesn't look like the one on the slide.

Teaching Tips

- You may want to describe some of the devices in a little more detail. For example, you may want to have a cartridge, DAT and reel tape available in the class so students can see the difference. It is also very helpful to point out the different devices in the classroom as you go through this.

Transition

In the upcoming sections of this module we will trace the I/O path from the SPU to the I/O interfaces (I/O cards and cables), and finally to the individual I/O devices (or peripherals).

Before we look at the actual SPU's, let's look at the types of interfaces found in the SPU for connecting devices.

2-3. SLIDE: Peripheral Interfaces

Peripheral Interfaces

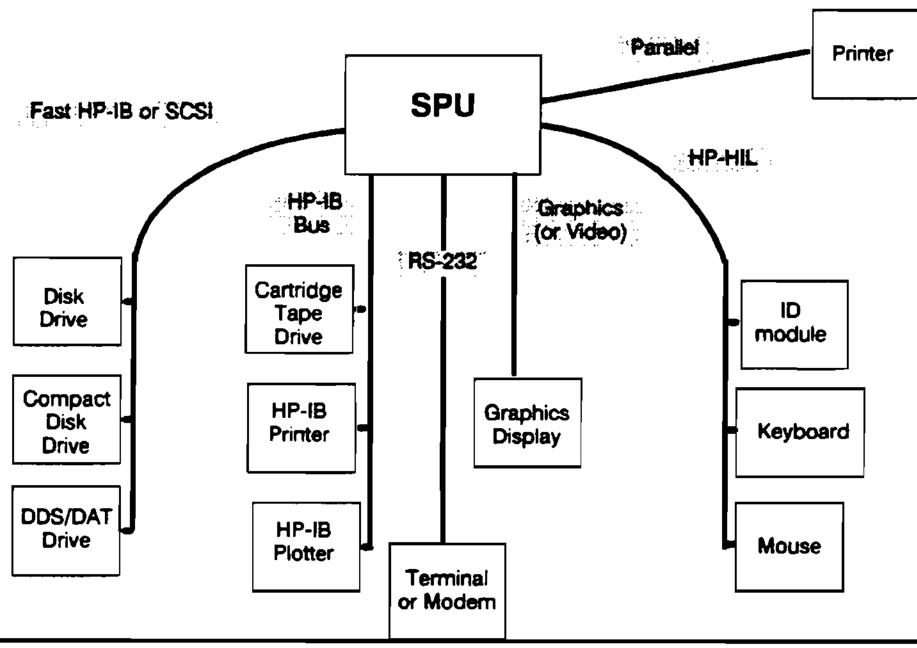

50722 2-3 9 ©1992 Hewlett-Packard Co.

Student Notes

There are many different types of interfaces supported on the system. It is necessary to connect the devices to the correct interface using the correct interface cable.

Interface Types

HP-IB — Hewlett-Packard Interface Bus. Originated by HP to allow electronic instrumentation to talk to computers. Supports disk drives, tape drives, CD-ROM drives, printers, plotters, and miscellaneous instrumentation. This bus may also be referred to as the IEEE-488 bus or GPIB (general purpose interface bus).

SCSI — Small Computer Systems Interface (pronounced scuzzy). A popular interface for connecting disk and tape drives to personal computers. It's capabilities have been improved and it is now finding expanded use in other parts of the computer marketplace. Supported devices include disk drives, tape drives, and CD-ROM drives. There is a second type of SCSI called **Differential SCSI**. This is a faster version of the SCSI standard to support faster devices such as disk arrays. *Differential SCSI devices and standard SCSI devices must* **NOT** *be connected together on the same interface card.*

HP-FL or PBA-FL — Hewlett-Packard or Precision Bus Fiber Link. A high performance proprietary fiber optic disk interface (currently available only to Series 800 systems).

Serial — The standard RS-232 device interface. This interface is used for connection to a terminal, modem, printer, plotter, or one of a number of various other devices. Many systems have one or more built in serial interfaces each supporting only one device. Some systems support multiplexers which provide panels containing four to sixteen serial interfaces (called **ports**).

Parallel — Centronics Parallel printer interface. An interface used mostly for connecting high speed printers such as LaserJets. This interface is faster than serial interfaces so if you have a choice of printer interfaces, parallel is usually more desirable.

HP-HIL — Hewlett-Packard Human Interface Link. This is a proprietary bus used to connect human input devices associated with graphics displays in a daisy-chain manner. Device types consist of keyboard, mouse, track-ball, knob box, digitizing tablet, button box, touch screen, ID module, and other third-party input devices.

LAN — Local Area Network for IEEE-802.3 or Ethernet. Almost all systems are equipped with either a BNC thin-LAN connector or a 15-pin Attachment Unit Interface (AUI) connector used to connect to thick-LAN or Ether-twist networks. Many systems have both and you must change an internal jumper to select which one will be active.

Video — Graphics Video Interface. This interface will only be found on systems that use graphics displays. On most models this interface consists of 3 BNC style connectors labeled R (red), G (green), and B (blue) which will allow you to connect a color graphics display using an RGB cable. Other graphics interfaces connect to another box called a graphics accelerator. This separate accelerator box will have the RGB connection for the display.

Audio — Audio input/output. Many workstation models (3/4/700) supply a plug for connecting external "beepers". Some models also provide a plug for connecting audio input devices.

2-3. SLIDE: Peripheral Interfaces

Instructor Notes

Purpose

Our intent here is to introduce the variety of interface types and describe briefly what each does. We don't want to blow the students away with the number of interface types so try to explain that the number is due to the dynamic nature of the computer industry at this point in time.

Teaching Tips

- Point out examples of the various cables and connectors on your lab systems and give the students time to peek and poke around a bit. This typically makes a good exercise at lab time.
- You may also want to bring up HP-FL which does not show up on the slide. It only supports fiber optic disks.

Transition

There are many things to know about all the different types of interfaces. Let's take a look at each one separately.

2-4. SLIDE: SCSI Devices

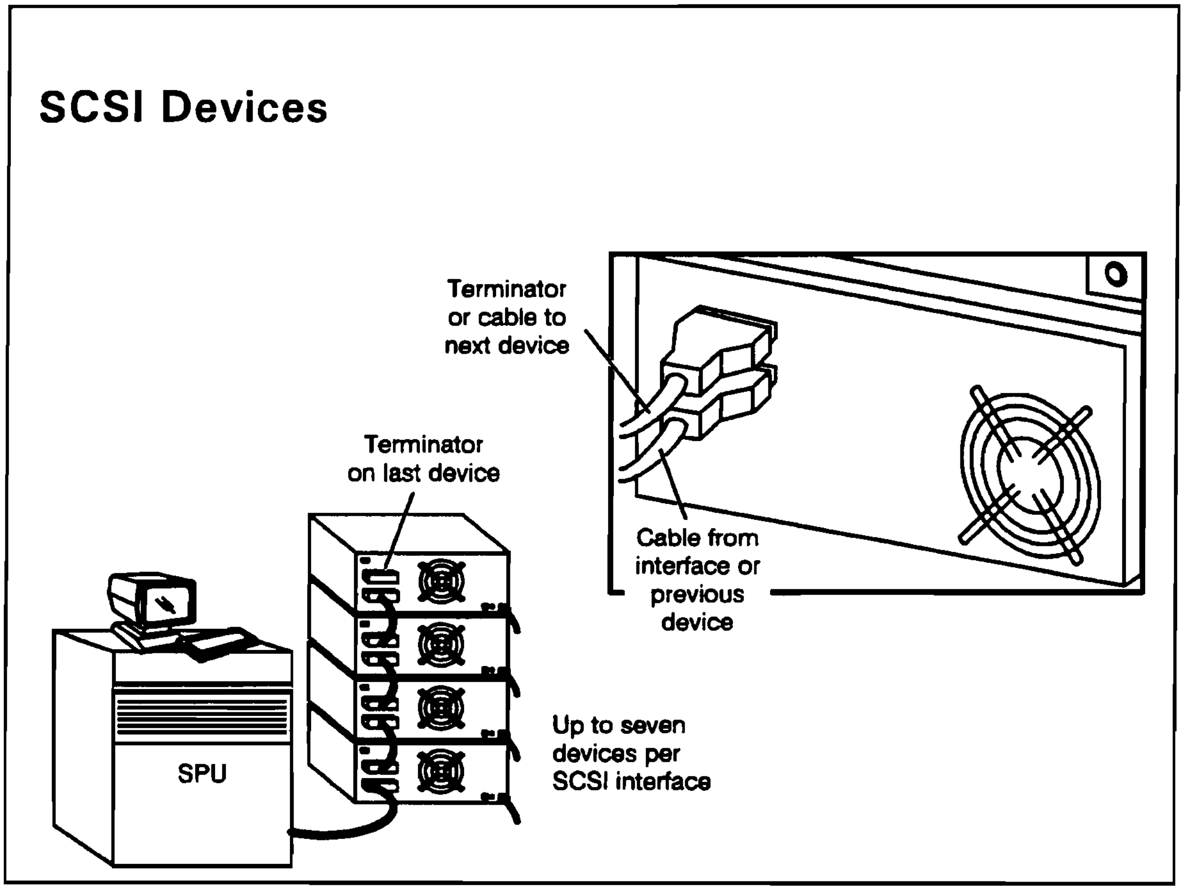

50722 2-4 10

Student Notes

Most SCSI devices are mass storage peripherals. These devices include disk drives (hard disks, magneto-optical, micro-floppy and compact disks) and tape drives (1/2 inch mag tape and DDS/DAT). SCSI interface cards support multiple SCSI devices using daisy-chain type connections between devices as shown on the slide. Every SCSI device has two ports on the back to facilitate the connection from the previous device to the next device.

Definition

If an SPU comes standard with a SCSI interface it is a **Single-Ended SCSI-II** interface, which is also known as normal or standard SCSI-II or just SCSI.

The **Differential SCSI-II** interface is a faster version of SCSI that plugs into a Series 700 EISA slot.

Here are some general guidelines for using SCSI devices with your system.

- Single-ended and differential SCSI II devices **MUST NOT** be mixed on the same interface.
- **DO NOT** connect or disconnect any SCSI device while the system is running, or turn power on or off to any SCSI device while connected to a powered-up system. This could result in corrupt data on the SCSI bus.
- Keep all devices powered on during and after system boot-up.
- Do not add or remove SCSI devices while the system is powered on.
- A maximum of **seven** devices are supported on a single SCSI interface card. All devices must have an **unique bus address** between 0 and 6. This address is set on the back of each SCSI device.
- If your **system disk** (the one containing the HP-UX operating system) is a SCSI disk it should be at **address 6**.
- Ensure that the total **cable length** (including external and internal cables) does not exceed 6 meters. The length of the SCSI bus should be kept as short as possible. However, do not use cables less than 0.5m in length.
  - Differential SCSI can have a total cable length of 25 meters.
- The last SCSI device in the chain, even if it is the only one, must have a **terminator** installed to its second connector. Without the terminator no device on the bus will work.
- If you have no external devices you must terminate the SPU connector to ensure the internal devices work properly.
- There are two types of terminators high-density (smaller) and low-density (larger). High-density terminators typically fit the connector on the SPU and low-density terminators fit the connectors on devices.

2-4. SLIDE: SCSI Devices

Instructor Notes

Purpose

SCSI is becoming the standard disk interface on all HP systems. It is very important the students know a little about using it.

Key Points

- Don't forget to cover SCSI terminators.
- Bus addresses are set on the back of each device using one of a variety of switch mechanisms. This will be discussed soon.
- If students have one, make a note that the C1700A, Magneto Optical Disk Autochanger, requires 3 (three) SCSI addresses (one for the autochanger and one for each of its two drives).

Teaching Tips

- Show the different types of SCSI cables terminators and how devices are connected.

Preparation Notes

There are many SCSI cable configurations. The *HP Apollo 9000 Series 400 HP-UX Owners Guide* under *SCSI Cables for Series 400 Workstations and Their Peripherals* will give you a good idea of the different types of cables used on all systems..

The *Pricing and/or Configuration Guides* (available from the Literature Distribution Center) also have good device and cable reference information. The guides for different SPU types will also help you answer questions regarding which devices are supported on which SPUs.

Another excellent reference for SCSI guidelines is the *SCSI Device Guidelines* section in Chapter 1 of *Installing Peripherals Volume 1* (for all systems).

There are also detailed pictures of SCSI connections in all of the *Installing Peripherals* books every time a SCSI device is described.

Transition

The other popular interface for mass storage and other I/O devices is HP-IB.

2-5. SLIDE: HP-IB Devices

HP-IB Devices

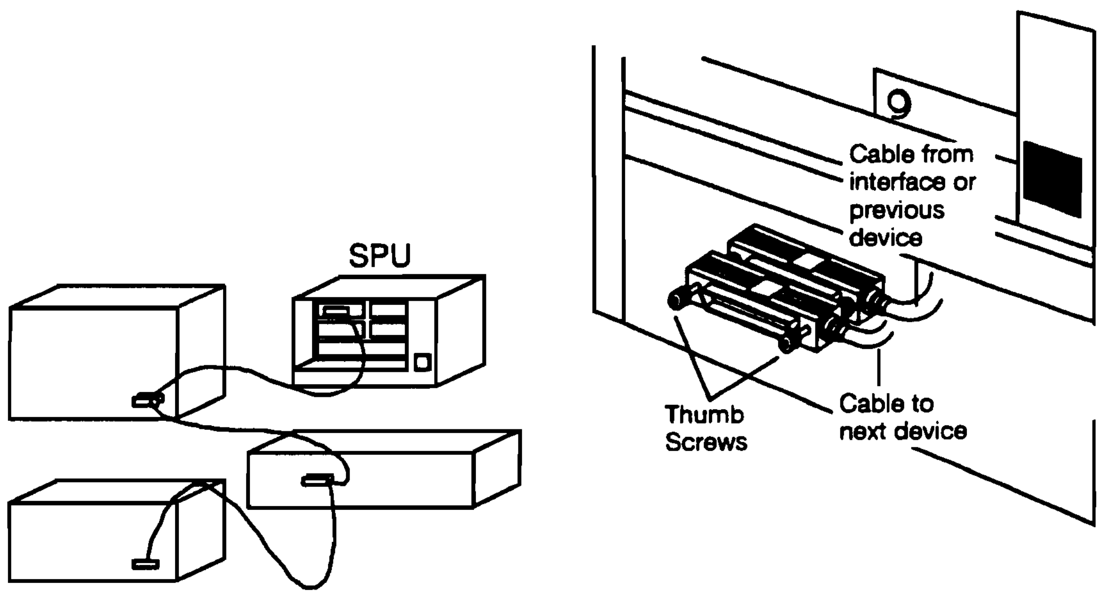

Up to 8 disks or 14 non-disk devices per HP-IB interface

50722 2-5 11

Student Notes

HP-IB is a general purpose interface used for connecting mass storage devices (various disk and tape drives) as well as printers, plotters, instruments and many other HP-IB (or IEEE-488) compatible devices.

HP-IB interface cards support multiple devices using the dual connectors on the HP-IB cables as shown on the slide. Every HP-IB device has only one HP-IB port. The cable, however, has two sides on each end. One side of this connector connects to the device then a second cable can connect to the back of that connector to connect to another device, etc.

There are two types of HP-IB interfaces:

Definition

Standard HP-IB (or just HP-IB) is used for any HP-IB type device. This interface is found on many Series 300, 400 and 800 SPUs.

The **High-Speed HP-IB** is used with faster mass storage devices, especially disks. This is a common interface to find on older Series 300 systems that do not have a SCSI interface.

The following list provides some guidelines for HP-IB devices:

- **DO NOT connect or disconnect** an HP-IB device while the system is running, or turn power on or off an HP-IB device while connected to a powered-up system. This could corrupt data on the HP-IB bus.
- A maximum of **fifteen non-disk devices** are supported on a standard HP-IB interface card. All devices must have an **unique bus address** between 0 and 14. This address is set on the back of each HP-IB device.
- A maximum of **eight disk devices** can be connected to an HP-IB interface. All devices must have an unique bus address between 0 and 7.
- If your **system disk** (the one containing the HP-UX operating system) is an HP-IB disk it should be at **address 0**.
- Slower devices (like HP-IB printers, cartridge tape drives, flexible disk drives, etc.) should be connected to the standard HP-IB interface.
- Faster devices (like disks) should be on the fast HP-IB disk interface (if your system is configured with one). Mixing fast and slow devices on the same HP-IB interface will reduce the performance of the faster devices.
- When only standard speed devices are used, the total HP-IB cabling on a standard speed interface is limited to 2 meters per device or 20 meters total, whichever is less.
- When only high speed devices are used, the total HP-IB cabling on a standard high interface is limited to 1 meter per device or 15 meters total, whichever is less.
- It is recommended that you do not stack the HP-IB cables more than 2 deep on any one HP-IB connector.

2-5. SLIDE: HP-IB Devices

Instructor Notes

Purpose

Key Points

- HP-IB devices do not require a terminator on the last device.
- Bus addresses are set on the back of each device using one of a variety of switch mechanisms. This will be discussed soon.
- A rule of thumb when mixing faster and slower HP-IB devices: the HP-IB bus is only as fast as the slowest active device on the bus. This means that having a printer and disk on the same interface will severely reduce the disk performance when the printer is active.
- Note that on systems with only one HP-IB interface there may be no choice but to mix devices and gamble with the performance.
- HP-IB works better in a bus configuration rather than a star. This is part of the reason for recommending that cables be stacked only 2 deep. The other reason is that the weight of more than 2 cables hanging off one connector may damage the interface or the screws on the cable.

Teaching Tips

- If possible, show how HP-IB cables work. Stress that you can't hook them up wrong.

Preparation Notes

See *HP-IB Device Guidelines* in Chapter 1 of *Installing Peripherals Volume 1* for more HP-IB information.

Although the HP-IB bus supports 15 devices by design, the CS80 (HP-IB disk) driver enforces a limit of eight open devices. You can, however, have as many as sixteen units for those eight devices. This allows devices like flexible drives with multiple units per bus address to avoid the limit of eight devices. If you connect more than eight disk devices to one interface, only eight can be used at a time.

Transition

Both SCSI and HP-IB devices must have unique bus addresses. We need to know how to set these addresses.

2-6. SLIDE: Setting Device Addresses

Setting Device Addresses

Device identified by bus address.

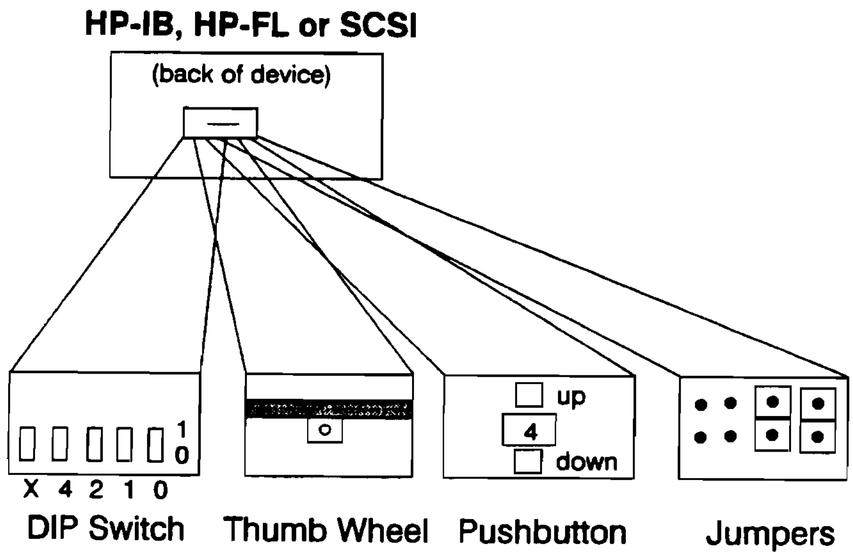

50722 2-6 12

Student Notes

If more than one device is physically connected to the same interface, the devices must be distinguished from one another by a device address (sometimes called the bus address).

On the back of each SCSI, HP-IB, HP-FL or PBA-FL device, you will find a mechanism for setting the devices's bus address. Common mechanisms include DIP (Dual Inline Package), thumb wheel, or push button switches. Note that devices internal to the SPU cabinet may use jumpers or DIP switches to set the device address.

For SCSI disks, it is suggested that the system disk be at address six (6). For HP-IB, HP-FL, or PBA_FL devices, the system disk should be at address zero (0). These choices are recommended due to the boot time disk search sequence and interface priority levels.

2-6. SLIDE: Setting Device Addresses — Instructor Notes

Purpose

Introduce the students to the various methods used for setting the individual device address.

Key Points

In most cases the main concern is that the address selected is unique for the bus that the device is connected to. This is a good time to let the students know that models with internal devices (eg. a 400 with an internal 210 Mb Rodime drive) will have pre-set bus addresses. Any external devices they connect must not conflict with the internal device addresses.

Teaching Tips

- If possible demonstrate on any available hardware where and how the switches are set.
- Note that dip switches are especially difficult because you don't know if you should read the switches from left to right or if up is on or off. The numbers on the DIP switch itself are not usually correct. You either have to look for an additional label or mark near the DIP or check the *Installing Peripherals* book or the device owner's guide for the devices for correct settings.

Transition

We sometimes need to connect serial devices to the SPU. There are also some guidelines you must follow when using serial devices.

2-7. SLIDE: Serial Devices

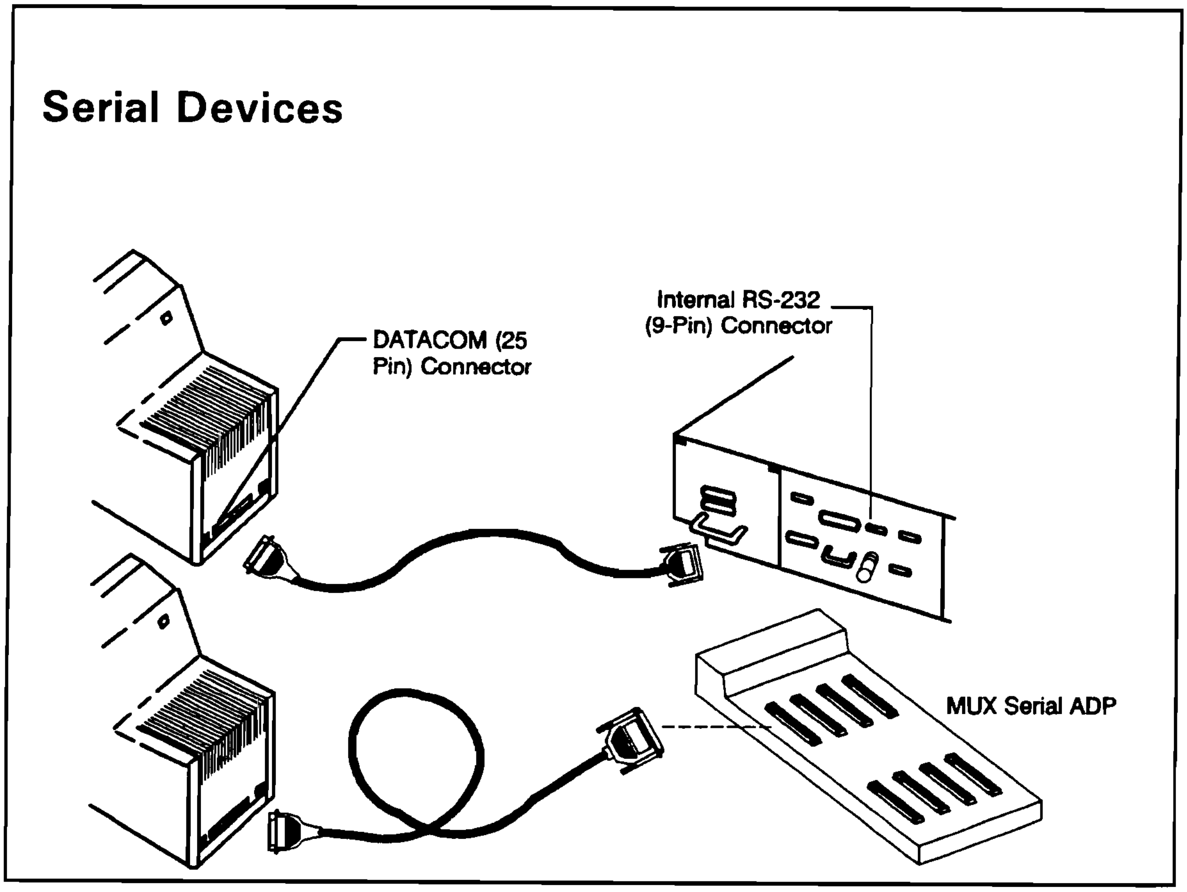

50722 2-7 13

Student Notes

Serial devices commonly include printers, plotters, modems and alphanumeric terminals.

One serial (RS-232) device is connected via a serial cable to a port that is either on the back panel of the system (an internal serial port) or a port on a distribution panel which is used with a multiplexer (MUX). Newer MUXs use an Active Distribution Panel (ADP) containing 8 or 16 serial ports. Both types of connections are shown on the slide.

There are many types of serial interfaces and cables. They vary in number of pins (usually 9 or 25) and gender (male or female). You must match the gender, proper number of pins and orientation of pins in order for the device to operate. This is the hardest part of configuring serial devices.

Reference

To find the exact cable needed to connect a device to your system, check the *Installing Peripherals* manual for your SPU type.

The owner's guide for the device should also help you.

If you cannot find the right cable for your serial device and SPU, contact your HP representative.

Serial devices are known to the system by what interface they are attached to. If the device is on an ADP, it is known by the **port number** where it is attached. This port number is printed next to the connector on the panel.

2-7. SLIDE: Serial Devices

Instructor Notes

Purpose

To show the different types of serial interfaces: internal and MUX with a multi-port distribution panel.

Key Points

- Newer systems (8X2s and 8X7s) use an Active Distribution Panel (ADP) with a MUX to offload some of the character processing from the actual MUX.
- On Series 800s where the system console is connected to an ADP, the first port (port 0) of the first ADP is the console port. On the 8X7 systems, port 7 is for the remote support modem (which can become the remote system console). Other Series 800s require a special ADP to connect the remote support modem. Check the owner's guides for details.
- Almost all serial interfaces are covered by the description in the notes. There is, however, a DIO 4 port mux that can be used on Series 300 or 400 systems. The 4 port mux has a 25 pin port (port 0) that can support modems, and 3 RJ-11 (phone jack type) ports that use a special cable supplied with the mux.

Teaching Tips

- Showing an ADP (if you have one) and having the students examine the internal serial ports and cables will help a great deal.

Preparation Notes

The *Installing Peripherals* manual has the numerous cable choices for Series 3/400 and 800 systems. The pricing and configuration guides have more complete information. You may want to look up some of the current part numbers for cables because the list is far to large to cover here.

Transition

The last type of device that we will discuss is parallel.

2-8. SLIDE: Parallel Devices

Parallel Devices

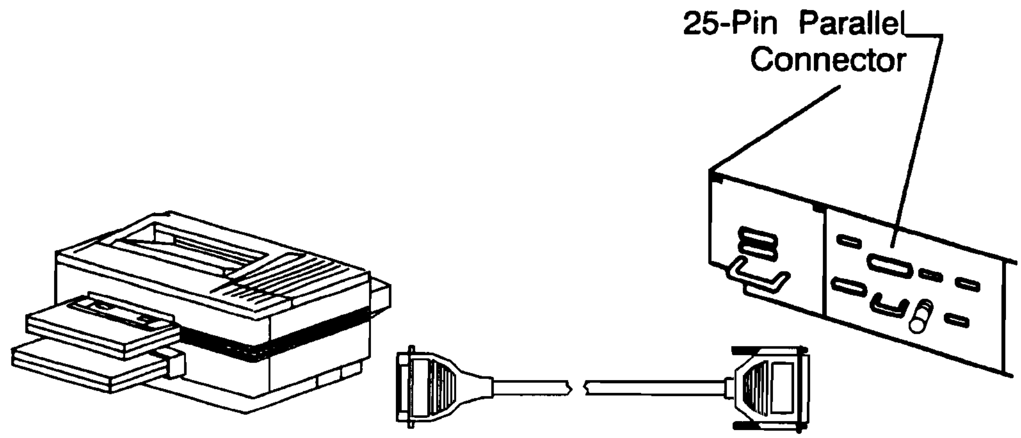

50722 2-8 14

Student Notes

The parallel interface (sometimes called Centronics) is used to connect a single devices such as a printer or image scanner.

This is probably the most straight forward interface to connect to and use. The cable has a 25-pin connector on one end to connect to the SPU and a large 36-pin plug on the other end that connects to the device. The only trick to the parallel cable is that the 25-pin end looks like, and will fit into a 25-pin serial interface. Make sure the interface you are connecting to is labeled "Parallel" and you will have no trouble.

Many devices have both a serial and parallel interface and you have to configure the device to use one or the other. This may be done using front panel controls (like the HP LaserJet II and III models) or using DIP switches (used on most other devices).

Reference

Check your device's owner's guide or the *Installing Peripherals* for your SPU to determine how to configure your device.

2-8. SLIDE: Parallel Devices

Instructor Notes

Purpose

We mentioned every other type of connection, we can't leave out parallel!

Key Points

- A common mistake is to plug the 25-pin parallel cable into the serial interface on the SPU. This is easy to do on the Series 400s because they have a 25-pin serial interface near the parallel interface. This will not work and could possibly damage the interface or device.

Teaching Tips

- Show the students a parallel cable and how it connects.
- This would also be a good time to show students how to set a LaserJet's interface to Serial or Parallel. This is in the *LaserJet User's Guide* but the steps are simple for a LaserJet II or III:

1. Take the printer off line by pressing the "ON LINE" button.
2. Hold the "MENU" button until the display changes to "Auto Cont=off" (II-D and III models) or "Sym Set=Roman-8" (II models).
3. Click the "MENU" button until "I/O Serial" or "I/O Parallel" appears in the display
4. Press "+" or "-" until the proper interface type appears
5. Press "ENTER"
6. Put the printer back on line by pressing the "ON LINE" button.

Preparation Notes

You may want to find the most current part numbers for cables in the configuration and price guides. The currently recommended parallel cable part number is 92284A (2.13m long). The 24542D cable used with PCs also works.

Transition

So much for background information on device types and interfaces. Now we have to discuss how we get HP-UX to recognize all of these devices we have been talking about.

2-9. SLIDE: How do I get there from here ?

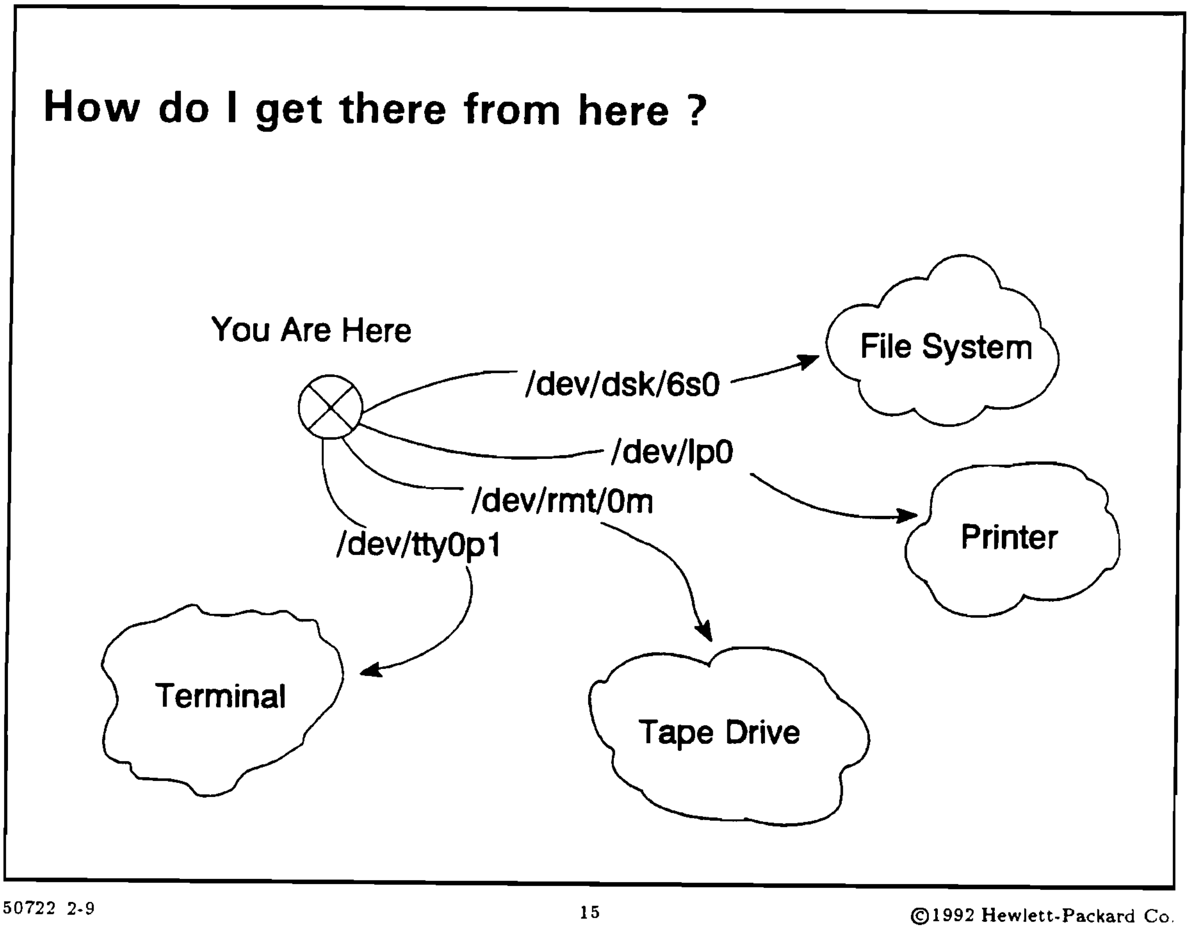

Student Notes

The HP-UX operating system has two primary responsibilities, they are to manage all the processes active on the system and to handle all the I/O requests between the active processes and the available I/O devices.

Definition

HP-UX references all of its devices as special files called **Device Files**. Doing I/O to a device then is treated the same as reading from and writing to regular files.

Each device file tells HP-UX two important pieces of information; which **device driver** to use and where the device is located on the system.

Definition

A **Device Driver** is a piece of software that must be part of the operating system kernel for the system to talk to a particular type of device. This device driver is often referred to as the device's **major number**.

The location and other characteristics about a device are collectively referred to as the **minor number**.

Most device files you need are already created for you. If you use SAM to add a new device to your system, it will make the appropriate device files. We will cover this in more detail when we talk about adding the different types of devices.

Reference

There is an excellent introduction to device files in the *Installing Peripherals* manuals in Chapter 1 of Volume 1 for Series 300/400/700 systems and Chapter 6 of Volume 2 for Series 800 systems.

LSDEV → MAJOR NO

ll → MINOR & MAJOR.

lssf → DECODES.

2-9. SLIDE: How do I get there from here ? Instructor Notes

Purpose

At this point we want to introduce the concept of device files and position them and their role in the flow of system I/O.

Key Points

- The concept of "device files" allows HP-UX to take a very simple approach to the interactive and programmatic handling of I/O. The basic premise is that all I/O is handled just like regular file I/O.

Teaching Tips

- Point out that "device files" are created with a naming convention that allows much easier access to peripherals than remembering cryptic addressing schemes.
- If students are logged in they can list the contents of the `/dev` directory to see some of the device files.

Transition

Device files are commonly found in the /dev directory.

2-10. SLIDE: Device Files in /dev

Device Files in /dev

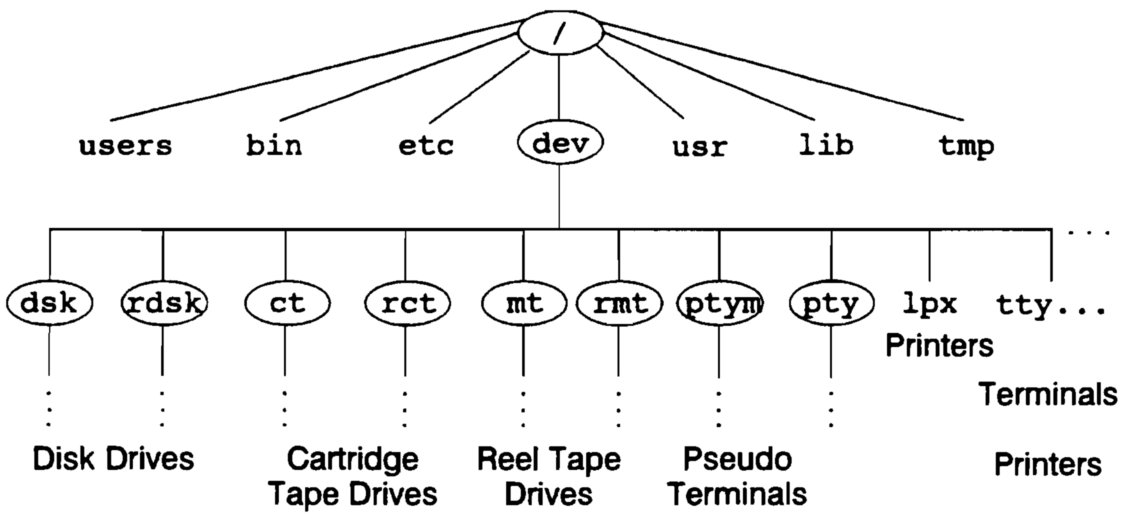

50722 2-10 16

Student Notes

This slide shows some of the default names of device files and subdirectories used by HP-UX. Most UNIX systems will have their device files in the /dev directory. The subdirectory names shown are for some of the different device types such as disk or tapes. Your system may not have all of the subdirectories and device files listed here.

Definition

Block devices contain an HP-UX file system and are referenced using block device files. This is how we access most disks.

Everything else is a **character** device. We reference these devices (like tape drives and printers) using character device files.

Following a directory and device file naming convention makes it much easier to locate device files for particular devices. The following table lists the device file naming conventions used by many tools in HP-UX (like SAM). You may use different device names but you will have to remember these names to access your devices.

Table 2-1. Device File Naming Conventions

Type of Device	Type of Device File	Naming Conventions Series 300/400/700	Naming Conventions Series 800
Disk	Block Character	/dev/dsk/*N*s0 /dev/rdsk/*N*s0	/dev/dsk/c*L*d*D*s*S* /dev/rdsk/c*L*d*D*s*S*
Cartridge Tape Drive	Block Character	/dev/ct/c*N* /dev/rct/c*N*	/dev/ct/c*L*d*D*s2 /dev/rct/c*L*d*D*s2
DAT and MAG Tape	Block Character	/dev/mt/*N*[l|m|h][n] /dev/rmt/*N*[l|m|h][n]	/dev/mt/*L*[l|m|h][n] /dev/rmt/*L*[l|m|h][n]
Printer or Plotter	Character	/dev/lp_*printername*	/dev/[r]lp*L* /dev/lp_*printername* /dev/tty*L*p*P*
Terminal	Character	/dev/tty*PP*	/dev/tty*L*p*P*
Modem, dial-in Modem, dial-out	Character	/dev/ttyd*PP* /dev/cua*PP*,/dev/cul*PP*	/dev/ttyd*L*p*P* /dev/cua*L*p*P*,/dev/cul*L*p*P*

Legend:

N—This is typically the HP-IB, SCSI, or HP-FL device address.

L—Series 800 only. The Logical Unit (LU) number assigned to the device on boot up. The first device of a particular type (printer, disk, MUX, etc.) is LU 0, the next is LU 1, and so on.

PP or *P*—The serial interface number or port number (if the device is on a MUX).

D—The Device Unit number of a peripheral (almost always 0).

S—Series 800 only. The Section number of a hard disk drive. (on series 800 machines a physical disk drive can be divided into several separate sections, each is treated as a separate device).

[l|m|h]—On mag tape drives this indicates the recording density in bytes per inch (l=800 bpi, m=1600 bpi, and h=6250 bpi).

[n]—The optional n shows that the device file has been configured so the tape will not rewind when you are finished accessing it.

tty*—Terminal device files. The name stands for teletype.

cul/cua—These two device files are used for dial-out modem operation. They stand for call-unit-autodialer (the device file used to dial the modem) and call-unit-line (the device file to communicate with the other system).

2-10. SLIDE: Device Files in /dev

Instructor Notes

Purpose

To introduce the naming conventions for the differnt devices on the different SPUs.

Key Points

- One of our primary goals is device file identification so the students can effectively find and use devices on their systems. This slide is very important to that objective.
- Most device files will be created by SAM or auto-configured. There are rare occasions when the system administrator will have to create device files manually.
- The Series 800 automatically configures in and makes device files for any new peripheral that is attached at boot time. Students should not have to get into making device files at all unless they want to rearrange their device connections or bus addresses on the system. We can only hope that the workstations will do the same in the near future.
- The Series 700 is installed with almost every device file in place. The only time a device file will have to be made is if the existing device file references the wrong device. For example, `/dev/rmt/0m` is the name for a DAT device at SCSI bus address 3. If the DAT drive is at any address other than 3 the device file will have to be remade. You might want to point out that on a Series 700 it may actually be easier to change the device address of the DAT drive in this case.
- Series 3/400 systems have relatively few device files installed. Most needed device files can be made by SAM. Some, like DDS/DAT device files cannot. SAM also "messes up" when creating the device file for a parallel printer. Fixes for these two problems will be given when we discuss identifying device files by major and minor number.
- One device, DAT, breaks the naming conventions on the Series 3/4/700s. This device is typically called `/dev/rmt/0m` regardless of its bus address. This is because many tools like SAM, frecover, and tar use this name as their default tape drive name. You may also want to mention that the density specifications are meaningless on a DAT drive.
- A second device breaks the naming conventions on Series 3/400. This is the root disk device file. It is always called /dev/dsk/0s0 regardless of its bus address. This can be very confusing if the root disk is SCSI and is really at address 6.

Teaching Tips

- Point out the naming convention table and go through some of the examples. You should explain how things like LU, port, and section numbers on the Series 800 work(if applicable). For example a terminal on port 1 of the second mux would have the device file name `/dev/tty1p1`.
- Point out device file names of important devices like the root device, tape drive, and printer. More examples will be given on the next slide.
- Device files like /dev/update.src and /dev/root should also probably be pointed out depending on the level and needs of the class.

- This would be a very good time for the students to look around in the /dev tree to see what exists on their systems.

Some Standard Device Files

This is a list of some of the device files that don't deal with the administrator configurable devices. These are in /dev and no doubt you will get questions on them:

Device File	Description
/dev/console	Physical system console (system message port)
/dev/syscon	Active system console (usually linked to console)
/dev/systty	System tty (always linked to console)
/dev/tty	Process group control terminal (not like /dev/ttyPP!)
/dev/null	Null file (system trash can or "bit bucket")
/dev/mem	Physical memory image
/dev/kmem	Kernel virtual memory image (the running kernel)
/dev/swap	Swap device(s)
/dev/[r]root	Character/Block root disk device
/dev/update.src	Character device used during system install
/dev/bsrc	Block device used during system install
/dev/hil*	HP-HIL devices (like the mouse, ID module, etc.
/dev/pty*	Master pseudo terminal devices (linked into /dev/ptym)
/dev/tty[p|q|r..]*	Slave pseudo terminal devices (linked into /dev/pty)

There are many others dealing with LAN, diagnostics, graphics, etc. that you typically don't have to worry about. Some clever keyword searches in *LaserROM* would help you find more. The *Installing Peripherals* will also help.

Transition

Let's take a look at some examples of device file names.

2-11. SLIDE: Name That Device

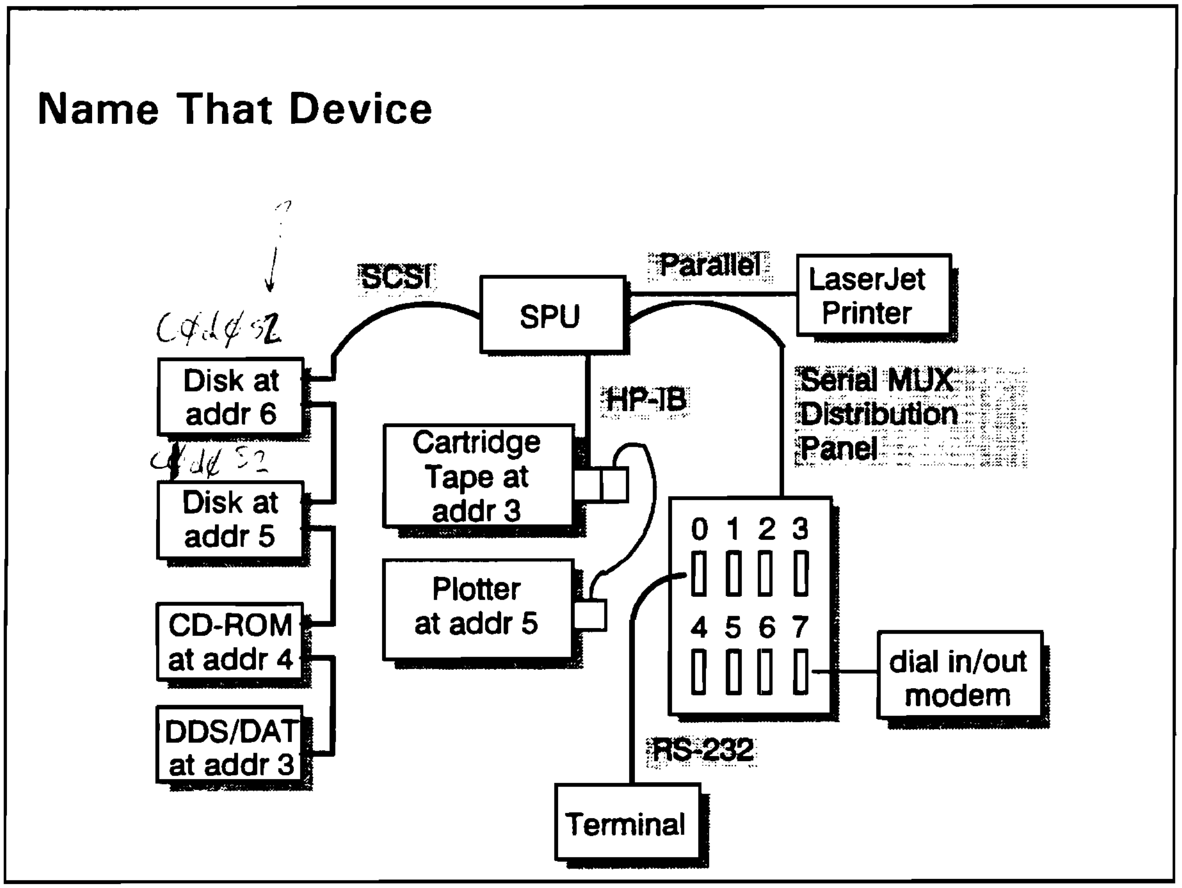

50722 2-11 17

Class Exercise

Working with the instructor and other members of the class, see if you can associate device file names with each of the devices on the slide.

You should use what you know about the various devices and interface types and the proper device file naming conventions as shown previously.

This is a fictitious system, of course. Few (if any) systems would have this combination of devices.

2-11. SLIDE: Name That Device

Instructor Notes

Purpose

To practice identifying devices and their associated device files.

Key Points

- Again, this exercise leads us to one of the main module objectives.

Teaching Tips

- If you can shine the overhead on a white board you should have plenty of room to write device file name examples next to the pictured devices.
- Be careful that the slide with your writing on it doesn't get too crowded. Using different colors for different types of devices will also help.
- If you have a mix of Series 800 and workstation customers you will have to be careful not to confuse the two different device file names for disks, cartridge tapes and terminals.
- You could also make copies of the slide to hand out so students could take separate notes for each type of machine that interests them.

Solution to Class Exercise

Table 2-2. Device File Names for Example System

Interface	Device	3/4/700 Name	800 Name
SCSI	Disk at addr 6	/dev/dsk/6s0 /dev/rdsk/6s0	/dev/dsk/c0d0s*S* /dev/rdsk/c0d0s*S*
	Disk at addr 5	/dev/dsk/5s0 /dev/rdsk/5s0	/dev/dsk/c1d0s*S* /dev/rdsk/c1d0s*S*
	CD-ROM at addr 4	/dev/dsk/4s0	/dev/dsk/c2d0s2
	DDS/DAT Tape at addr 3	/dev/rmt/0m	/dev/rmt/0m
HP-IB	Cartridge Tape at addr 3	/dev/rct/c3	/dev/rct/c3d0s2
	Plotter at addr 5	/dev/lp_*plottername*	/dev/lp_*plottername*
Parallel	LaserJet III	/dev/lp_*printername* /dev/prn_parallel	/dev/lp_*printername* /dev/lp0
Serial MUX DP	terminal on port 0	/dev/tty00	/dev/tty0p0
	dial-in modem on port 7	/dev/ttyd07	/dev/ttyd0p7
	dial-out modem on port 7	/dev/cul07 /dev/cua07	/dev/cul0p7 /dev/cua0p7

This would be a good time to remind students that a DDS/DAT drive is typically called /dev/rmt/0m regardless of its bus address.

Transition

Now that we know a little about the different types of interfaces on the systems we can look at some SPUs and how the interfaces are laid out.

2-12. SLIDE: Series 800 Model 827 SPU

Series 800 Model 827 SPU

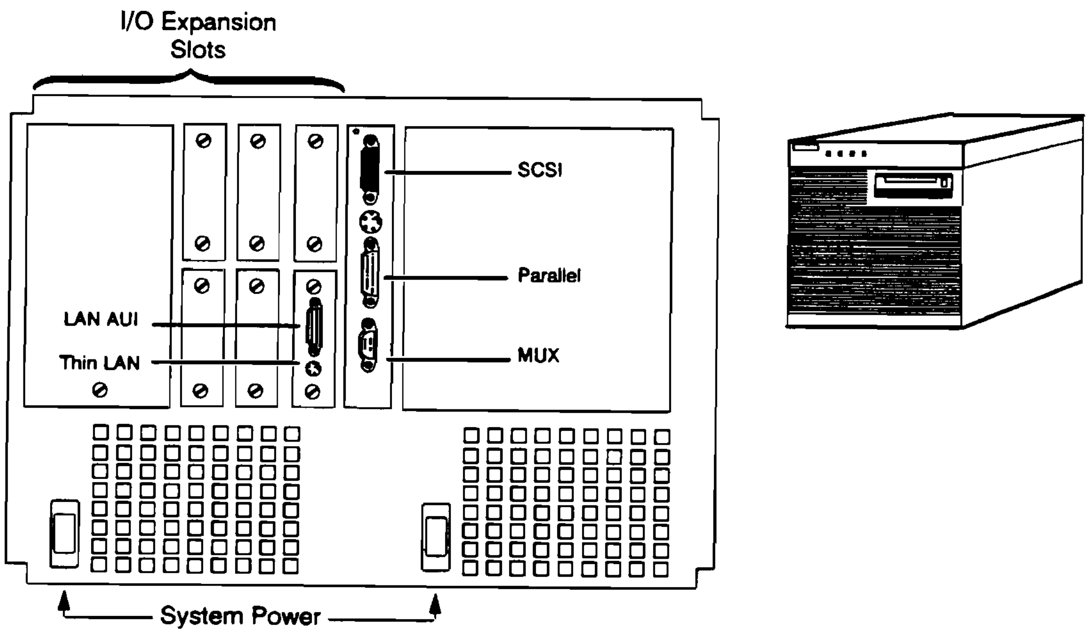

50722 2-12 18

Student Notes

When you are working with peripherals on an HP9000 Series 800 computer, you need to know a little about the Series 800 hardware addressing scheme.

There are two types of architecture associated with HP 9000 Series 800 computers, **CIO** and **HP-PB**. Dealing with peripherals on each of these architectures is similar, but differs somewhat due to differences in the layering of system busses.

Definition

CIO is *Channel Input/Output*. It is a proprietary I/O bus architecture which supports a variety of interface cards from HP. It is separate from the internal system bus used with the CPU and RAM. It allows great flexibility and expandability of the I/O subsystem for many Series 800 systems.

HP-PB is the *Hewlett-Packard Precision Bus*. This is a smaller and simpler bus architecture that is used as the internal system bus as well as the I/O subsystem bus. It supports numerous HP interface cards as well as RAM and other system cards.

The model 827 and all 8X7 and 8X2 machines use the HP-PB architecture. This architecture allows you to plug interface cards into numbered slots on the back of the SPU. These cards then have the connector needed to attach the cable that provides the external bus to the device(s). The slot number of the interface card becomes very important later when we discuss addressing peripherals and identifying device files. Some of these machines will support internal devices such as disk drives, CD-ROM disk drives and DDS/DAT tape drives as part of the SPU.

The other bus architecture, CIO, is used on some of the larger Series 800 models (850, 855, 860, 865, and 870) as well as some earlier models (825, 834, 835, 845). In this architecture the interface cards do not actually have connectors on the back of the machine as we see on the HP-PB systems. Most, instead, have cables coming out the back that end with the proper connector to connect to the device(s). Later in this chapter we will also see that the way we address peripherals on a CIO machine is somewhat more complicated that it is on HP-PB machines.

Regardless of the model or interface card, the connectors providing the interface types we discussed are basically the same.

Reference

If you are interested in learning more about the different architectures and the differences between models see the *Installing Peripherals* manual for *HP 9000 Series 600/800 Computers*. Another excellent discussion of the architectures and hardware addressing is given in the appendicies of *Installing and Updating HP-UX* for *Series 600/800 Computers*.

The *Series 800 Owner's Guide* that comes with the SPU contains information and pictures of the front panel display and the different interface and device options. It will describe the internal devices on the system where applicable. It also has good information about the basics of running your system.

You should contact your HP representative for specific Series 800 configuration information.

2-12. SLIDE: Series 800 Model 827 SPU — Instructor Notes

Purpose

To show the layout of a system that will most Series 800 students will be familiar with (the 8X7 family is very popular). We also want to briefly mention the other types of Series 800s and ensure students that no matter what type they have, the peripherals will connect the same way.

Key Points

- It should be stressed that this page and the next few are here as a basic introduction to where the interfaces are located on the machine. It is not our intent to show every model of Series 800 available and describe it's configuration options (you can't even do that in the 5 day admin class!).
- Even though peripherals connect in basically the same way on all models the hardware addressing scheme will be different between PB and CIO models (especially the larger 850 class).

Teaching Tips

- If you have students that have other Series 800 models you should lead them to the proper section of one or both of the referenced books so they can get a feel for what their machine looks like and where the interface cards will be.
- Any questions regarding specific customer configurations should be directed to the appropriate HP representative.

Preparation Notes

The referenced books are also very good tools for the instructor to help prepare for the course. If you are unfamiliar with any of the Series 800 architectures (old or new) you should read through Chapter 1 of *Installing Peripherals* and Appendicies A-E of *Installing and Updating HP-UX*.

Another good source of information for you is the *HP9000 Series 800 Business Servers Configuration Guide* (or something like that depending on when you read this). These books always contain great information on slot and I/O configurations on all Series 800 platforms.

Transition

Let's take a look at some of the specifics with identifying the different interfaces in the Series 800 SPU.

2-13. SLIDE: Series 800 Hardware Addressing

Series 800 Hardware Addressing

To view the system hardware configuration and all devices use `ioscan`:

```
# ioscan -u

Hardware
  Path     Description        Status
===================================
4          hpib               ok
4.2            disk           ok
8          tty                ok
16         lan                ok
52         scsi               ok
52.0           target         ok
52.0.0             tape_drive ok
52.6           target         ok
52.6.0             disk       ok
53         printer            ok
56         tty                ok
```

50722 2-13 19

Student Notes

On a Series 800 you can see your complete hardware configuration using the `ioscan` command. The information on the slide is typical `ioscan` output for an 8X7 SPU. Other Series 800 systems will have different `ioscan` output but the general concept is the same.

Definition

The **Hardware Path** identifies the physical location of the device connected the SPU.

The first digit is called the **module number**. This is number is (4 x SPU slot number). The second digit is the device number (bus address). The third digit is only used for SCSI devices and is always 0 (zero).

This **hardware path** description is for HP-PB systems. Other Series 800 systems will have a slightly different (and more complex) hardware address because of the more complex I/O subsystem.

Reference

See chapter 1 of *Installing Peripherals* for the *HP9000 Series 600/800 Computers* for complete details of the hardware addressing schemes of other Series 800 systems.

The hardware address is very important because it is how we tell tools like SAM what hardware we are working with.

The **Description** field describes what type of interface or devices is at the listed hardware path. Some common types include:

Description	**Device**
hpib	HP-IB interface.
scsi	SCSI interface.
printer	Parallel printer interface.
lan	Network interface.
tty	Serial multiplexer (MUX).
tape_drive	DDS/DAT or mag tape drive.
disk	Disk drive.

2-13. SLIDE: Series 800 Hardware Addressing — Instructor Notes

Purpose

We have to know a little about hardware addressing because SAM asks which interfaces or devices to configure using the hardware address. We also want the students to be able to use `lssf` to find which device files represent a particular device (see the second slide [Continued] of this topic).

Key Points

- SAM uses the `ioscan` utility to display choices of hardware to configure.

Teaching Tips

- If you have students with non-HP-PB systems you should point them to the proper section in *Installing Peripherals*. It is fruitless to explain all of the addressing details to all of the students.
- Make sure students realize that no matter what type of architecture they have `ioscan` will report the correct hardware addresses.
- Only cover specific addressing schemes if there is a general interest from most of the students. (Do not spend a lot of time here if you have a classroom full of workstation users and spend a lot more time here covering H/W addressing if you have mostly Series 800 students.)

Other Series 800 Hardware Addresses

- Models 825, 834, 835, 845: **M.C.D.U**
- Models 850, 855, 860, 865, and 870: **B/M.C.D.U**

Where **M** is the module number (4 * the channel adapter slot number), **C** is the CIO slot number of the interface card, **D** is the device address and **U** is the SCSI unit number (currently always 0).

B is the bus convertor number for the larger Series 800 systems. This will be 2 for any interface card in the right half of the system and 6 for any card in the left half of the system (as you face the back of the system).

See Chapter 1 of Volume 1 of *Installing Peripherals* for *HP 9000 Series 600/800 Computers* for all of the details.

Transition

First cover the second slide of this topic. The next Instructor Notes cover the transition.

2-13. SLIDE: Series 800 Hardware Addressing (Continued)

Series 800 Hardware Addressing (Continued)

Use lssf to identify device files:

```
# lssf /dev/dsk/c0d0s13
disc3 lu 0 section 13 address 52.6.0 /dev/dsk/c0d0s13
```

50722 2-13 20

Student Notes

Device files and hardware addresses

If you need to identify what device a device file references, you can use the `lssf` *device_file* command.

```
# lssf /dev/rmt/*
tape2 lu 0 bpi 6250 at&t address 52.0.0 /dev/rmt/0h
tape2 lu 0 bpi 6250 at&t compressed address 52.0.0 /dev/rmt/0hc
tape2 lu 0 bpi 6250 no_rewind at&t address 52.0.0 /dev/rmt/0hn
tape2 lu 0 bpi 800 at&t address 52.0.0 /dev/rmt/0l
tape2 lu 0 bpi 800 no_rewind at&t address 52.0.0 /dev/rmt/0ln
tape2 lu 0 bpi 1600 at&t address 52.0.0 /dev/rmt/0m
tape2 lu 0 bpi 1600 no_rewind at&t address 52.0.0 /dev/rmt/0mn
```

This gives you all of the information you need to relate any device file with a specific device on your system.

For the example of the 827 we have on the slide we can see that there is a SCSI disk at 52.6.0. This means that the SCSI interface is in slot 13 of the SPU and the disk is at bus address 6. Using `lssf`, we can find out that the device files for the disk are `/dev/dsk/c0d0s`*S* where *S* is the disk section number.

2-13. SLIDE: Series 800 Hardware Addressing (Continued)

Instructor Notes

See the prior Instructor Notes.

Transition

We have seen the basic setup of the multi-user Series 800s. The other major type of the HP 9000s is the workstation family. Let's take a look at their SPUs and interfaces.

2-14. SLIDE: Series 700 Model 720/730 SPU

Series 700 Model 720/730 SPU

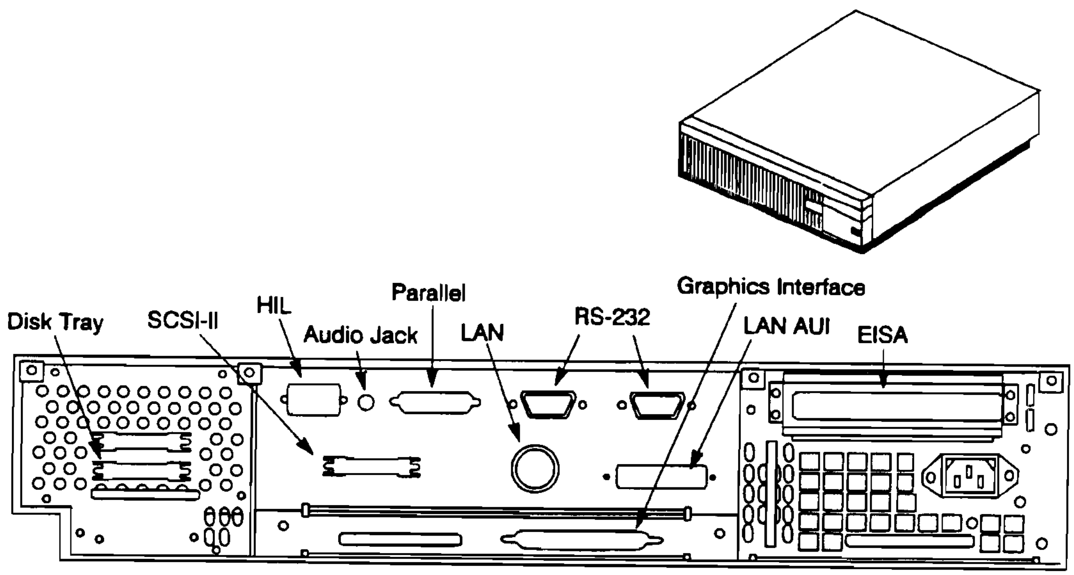

Rear Panel

50722 2-14 21

Student Notes

The Series 700 has most of the peripheral interfaces on its **Core I/O board**. Some models may also have one or more **EISA** slots for other interface cards like the Differential (fast) SCSI card.

Definition

EISA is the *Extended Industry Standard Architecture*. As the name implies it is a widely used bus architecture which supports a variety of interface cards from numerous vendors. EISA is used to expand the I/O system on Series 700 and some Series 400 systems.

The Series 700 models support a variety of internal devices including disk drives, DDS/DAT tape drives, CD-ROM drives and 3.5 inch micro-floppy disk drives. Not all models will support all types of devices.

On models 720 and 730 (pictured on the slide) you can see the core I/O board in the center of the machine. You may also have a separate video interface card below it for connecting the graphics display or a graphics accelerator. Notice that these models also have an internal removable disk tray that can hold two SCSI disks. This tray is connected to the SCSI interface on the Core I/O board using a short SCSI cable provided with the SPU. The 720 and 730 may each have one EISA slot.

The other desktop Series 700 models include the models 705 and 710. These are smaller SPUs and all of the interface connections are on the back of the unit. You can not distinguish the video interface from the Core I/O board. As with the other models the 710 may have up to two internal disks. Unlike the 720 and 730, however, these disks are not removable. The 705 and 710 do not have any EISA slots available.

The largest model in the Series 700 family is the 750. It is a desk side model that can also contain many internal devices (more than the desktop models). The biggest difference in the 750 is that it has four EISA slots for expansion and can support more than one video interface.

All models reference devices and device files in the same fashion. The major difference is the expandability of the SPU and the performance levels of the various models.

Reference

The *Series 700 Owner's Guide* that comes with the SPU contains pictures of the front and back panels, describes the different interfaces, and will describe the internal devices available on the system. It also has good information about the basics of running your system. There is also a hardware installation guide that comes with your SPU that has details on interface connections.

The *Installing Peripherals* for the *HP 9000 Series 700* contains very good and detailed information regarding the bus architecture and supported interfaces and peripherals.

You should refer to the *Pricing and Configuration Guide* or contact your HP representative for specific Series 700 configuration information.

2-14. SLIDE: Series 700 Model 720/730 SPU — Instructor Notes

Purpose

Again, this is a brief introduction to the Series 700 family.

Key Points

- The basic difference between models is performance, the actual SPU layout and EISA expandability.
- All peripherals connect to any model in the same way. Devices and device files are referenced the same regardless of SPU model.

Teaching Tips

- You may want to mention that we stay away from specifics here because the list of supported devices and system configurations changes too fast to keep up to date in a course.
- HP sales offices have the *Series 700 Pricing and Configuration Guide* that contains the most up to date information regarding supported peripherals and SPU configurations. You should consult this guide or have the customer contact their HP representative if they have questions regarding specific configurations.
- If you have a pricing and configuration guide in the classroom you may want to show it to students so they can look up model specifics. There is no good documentation provided to customers on this.

Preparation Notes

Reviewing the pricing and configuration guides (available from HP) and any available owner's manuals will help you see the configuration of the different models.

Chapter 1 of the *Installing Peripherals* manual discusses the specific hardware addressing schemes and has some good pictures of the internal bus architecture.

Transition

The Series 700 has a unique hardware addressing scheme to identify its interfaces and devices.

2-15. SLIDE: Series 700 Hardware Addressing

Series 700 Hardware Addressing

The Core I/O board is known as card slot 0.

Each interface on the Core I/O board is known by a function number:

Function Number	Interface
1	SCSI
2	LAN
3	HP-HIL
4	Serial Port 1
5	Serial Port 2
6	Parallel

Device files have a **minor number** to identify their device:

0x20*FDXX*

F = Function Number *D* = Device Address

50722 2-15 22

Student Notes

The Series 700 core I/O board is addressed as **card slot 0**. SAM will need this number to address any interface on the core I/O board. Using the `Help` facility in SAM will let you select the proper interface on the core I/O board.

Each interface on the core I/O board is identified by a **function number**. Function numbers are assigned as described on the slide.

Device files and hardware addresses

Using the `ll` command you can identify which device files reference the specific devices:

```
# ll /dev/dsk
total 0
brw-r-----   1 root     sys          7 0x201000 May  1 09:35 0s0
brw-r-----   1 root     sys          7 0x201100 May  1 09:35 1s0
brw-r-----   1 root     sys          7 0x201200 May  1 09:35 2s0
brw-r-----   1 root     sys          7 0x201300 Jun 22 18:14 3s0
brw-r-----   1 root     sys          7 0x201400 May  1 09:35 4s0
brw-r-----   1 root     sys          7 0x201500 Jun  4 15:28 5s0
brw-r-----   1 root     sys          7 0x201600 Jun 18 09:11 6s0
```

Definition

The number that starts with `0x` is the **minor number**. This number identifies the physical device that the device file references.

The general form of the minor number looks like this:

0x20*FDXX*

Where *F* is the function number of the interface as described above, *D* is the device address (used for SCSI) and *XX* is device specific information (on a serial device file, for example, this will tell you if the device is used for dial-in, dial-out or if it is directly connected). We are most interested in the function number and device address.

From the long listing above of `/dev/dsk` you can see that these are SCSI disks (function number 1), the bus addresses range from 0-6 and the device files follow the naming convention we have described.

Reference

Chapter 1 of *Installing Peripherals* for *HP 9000 Series 700 Computers* has a more detailed look of the addressing scheme. It also has information regarding the addressing scheme for EISA devices. They do not use the same minor number format as the core I/O devices.

2-15. SLIDE: Series 700 Hardware Addressing — Instructor Notes

Purpose

Once again we need this information for two reasons:

1. It is easier to understand what SAM is asking for if we know a little about how the interfaces are addressed. (Actually, SAM on the 700 requires you to know less about the physical hardware than any other SPU type).
2. One of our major objectives is to be able to find device files for a particular device.

Key Points

- The minor number format is different for EISA interfaces. If you have Series 700 students with EISA slots you should probably mention the EISA scheme:

  0x4*S*0*DXX*

  Where *S* is the EISA slot number, *D* is the device address and *XX* is device specific information.

Teaching Tips

- You may want to introduce the `dmesg` command to identify what hardware is installed in the system. This works best shortly after the system comes up. If the system has been running for some time you may have received system (especially SCSI) errors that scrolled the I/O configuration information off of the `dmesg` listing.

Transition

One of the other popular workstation lines is the Series 400.

2-16. SLIDE: Series 400 Model 425T SPU

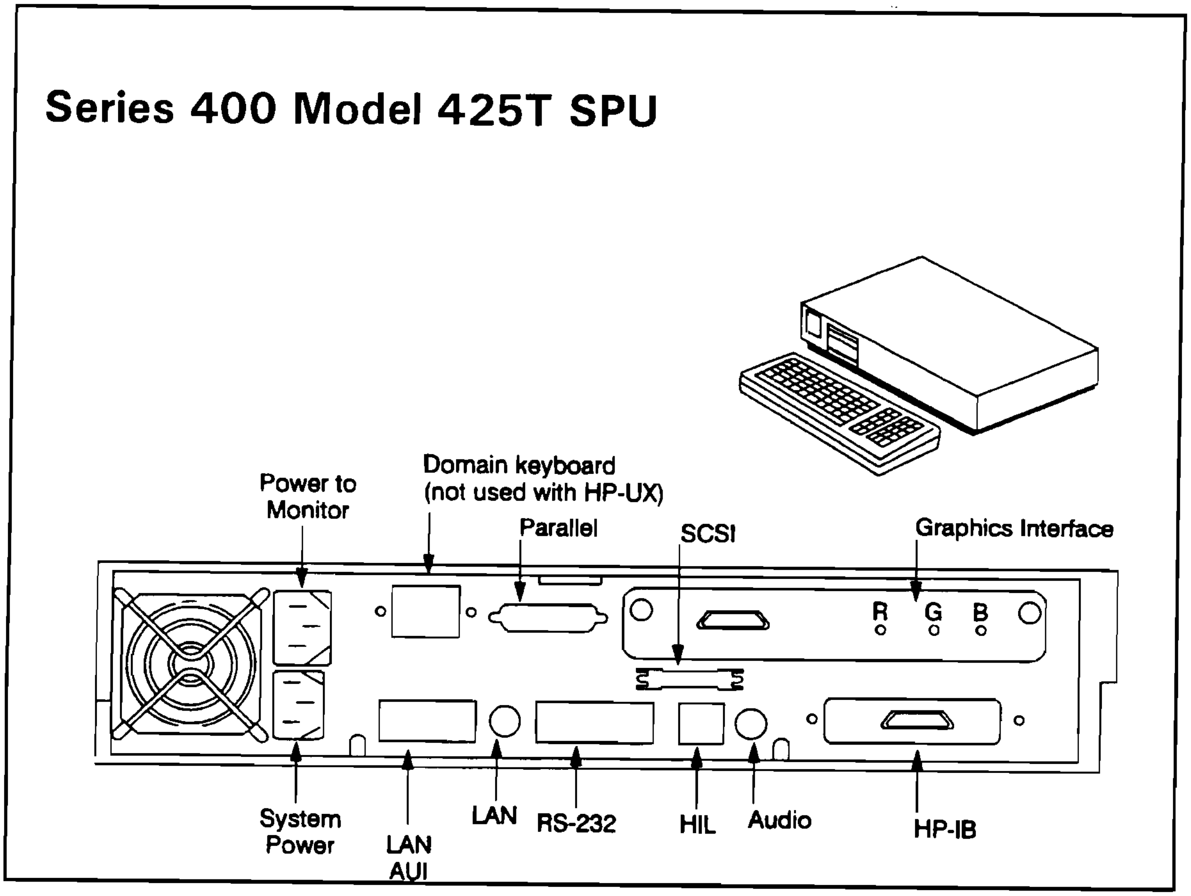

50722 2-16 23

Student Notes

The interface connectors for the Series 400 models appear directly on the back panel. All models support basically the same interfaces. The location and number of interfaces will vary.

As with the Series 800 and 700 the Series 400 systems support various internal devices including disk drives, DDS/DAT tape drives, CD-ROM drives and 3.5 inch micro-floppy disk drives. Not all models will support all types of devices.

The 425T ("T" for deskTop model) and the 400DL ("DL" for DiskLess) have the same SPU configuration as shown on the slide. The 400DL, being a diskless model will not have the internal disks or SCSI interface found on a 425T.

The "S" (for deskSide) models of the Series 400 (400S, 425S and 433S) are vertically oriented SPUs with expansion slots on the side. There are either three **DIO-II** slots or four **EISA** slots.

Definition

DIO (Direct Input/Output) is a bus that is used primarily for I/O interface cards and **DIO-II** is used for system cards such as RAM or video. A DIO-II slot can be change into two DIO slots using an adapter.

All Series 400 systems (except the 425E) have one DIO-II slot that is used for the video interface card. This video slot can not be used with the DIO adapter. The "S" models may have three additional DIO-II slots.

The 425E is a smaller SPU package where all interface connectors (including video) are fixed in the back of the SPU (no removable cards). The 425E does not have an HP-IB interface.

Reference

The *Series 400 Owner's Guide* that comes with the SPU contains information and pictures of the different interface and device options and will describe the internal devices on the system. It also has good information about the basics of running your system.

The *Installing Peripherals* for the *HP 9000 Series 300/400* contains more detailed information on connecting and accessing the supported interfaces and peripherals.

You should refer to the *Pricing and Configuration Guide* or contact your HP representative for specific Series 400 configuration information.

2-16. SLIDE: Series 400 Model 425T SPU — Instructor Notes

Purpose

Introduction to the Series 400 backplane layout

Key Points

- There are many variations on the Series 400 SPU. We don't want to talk about all of them here. Just enough to get the students started with I/O interfaces.

Teaching Tips

- The best pictures of these machines are in the color hardware slide set.

Preparation Notes

There is no good reference for the Series 400 architecture like there is for Series 800 or 700 systems. The current *Pricing and Configuration Guide* should help. The Jan'92 printing is not a very good reference for Series 400s (it's ok for 700s). Rumor has it that there is a better Series 400 guide due out soon.

The sales literature may also be a help. You may want to contact your literature coordinator for Series 400 sales and technical information guide.

The *Owner's Guide* is a good overview reference.

Transition

The last workstation architecture we want to mention is the Series 300.

2-17. SLIDE: Series 300 Model 382 SPU

Series 300 Model 382 SPU

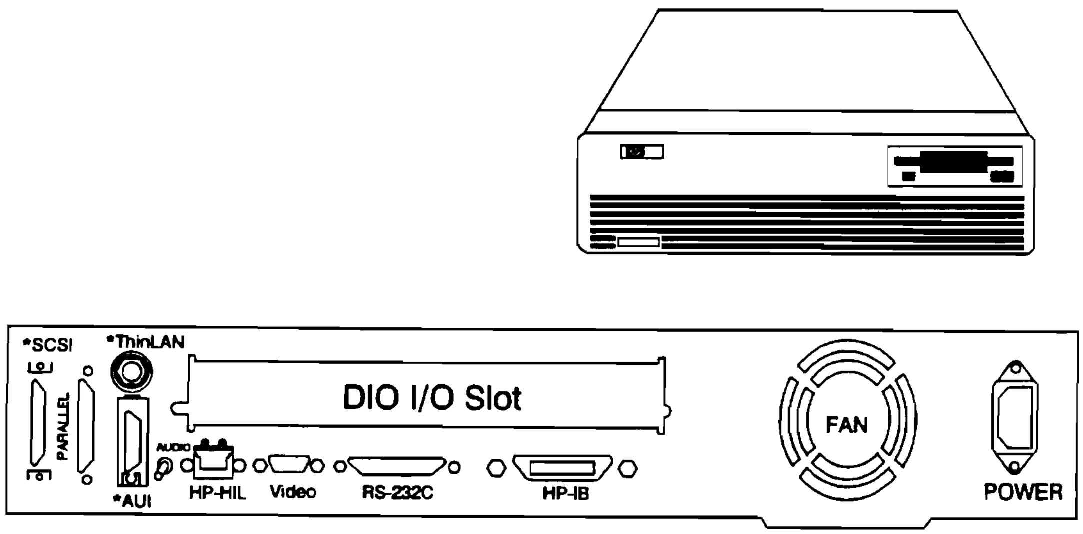

50722 2-17 24

Student Notes

The Series 300 is a general purpose workstation with great I/O capabilities. The classic Series 300s were simply four slot **DIO-II** card cages and you plugged in the CPU, RAM, **System Interface Card** and Video interface directly into the DIO-II bus. This is how the 330, 350,360, and 370 models are configured.

Definition

The **System Interface Card** is a single DIO-II card used on most classic Series 300 systems. It contained the basic I/O subsystem and interfaces including a disk interface (fast HP-IB or SCSI), LAN, serial, audio, keyboard and standard HP-IB interfaces.

On newer models the CPU, RAM and system interfaces are on the same card. The 345, 375, 380 and 382 are examples of this type of SPU.

Newer Series 300 SPUs like the 382 pictured on the slide, however, have their interfaces integrated directly onto the back of the SPU.

Most newer models of the Series 300 will also support one or more internal devices.

References

The *Installing Peripherals* for the *HP 9000 Series 300/400* has many pictures of the Series 300 models and detailed information about the various interfaces.

You should refer to the *HP 9000 Series 300 Family Pricing and Configuration Guide* or contact your HP representative to see the newest models and for specific configuration information.

2-17. SLIDE: Series 300 Model 382 SPU

Instructor Notes

Purpose

A very brief introduction to the Series 300.

Preparation Notes

The only models HP still sells that run HP-UX are the 380 and 382. It is unlikely that you will see many students with other types of systems. That should make your job easier.

The bad news is that there is very little information left around regarding the older Series 300s. Older configuration guides will help a great deal if you can find them.

Transition

We did not cover the Series 400 addressing scheme because it is the same as the Series 300. We will cover both now.

2-18. SLIDE: Series 300/400 Hardware Addressing

Series 300/400 Hardware Addressing

Series 300 and 400 systems identify interfaces by **Select Code**.

Select Code	Hex Notation	Interface
7	0x07	Standard HP-IB
9	0x09	Internal Serial
12	0x0c	Parallel
14	0x0e	Series 300 Disk Interface (HP-IB or SCSI)
14	0x0e	Series 400 SCSI
21	0x15	LAN

Device files have a **minor number** to identify their device: 0x*ScBaXX*

Sc = Two digit hex select code
Ba = Two digit device address

50722 2-18 25

Student Notes

The Series 300 and 400 systems use the same scheme for hardware addressing.

Definition

Interfaces on Series 300 and 400 SPUs are identified using **Select Codes**. This is a unique number set on each interface. The select code is changeable by the administrator but is typically left at the factory settings.

Select codes are supplied to the system in either decimal notation or in hexidecimal notation (base 16). Default select codes for the most common interfaces are listed on the slide.

You can verify the select code settings by running the dmesg command:

```
# dmesg

      .
      .

Internal HP-IB Interface - system controller at select code 7
  Parallel poll interrupts enabled.
HP98644 Advanced RS-232C Serial Interface at select code 9
  With 16 byte rcv fifo, 16 byte xmit fifo, hardware handshake, and high speed
clock.
Parallel Interface at select code 12
HP98265A (SCSI Interface) 5 MB/s; parity enabled at select code 14
      .
      .
```

Device files and hardware addresses

Like the Series 700, you must use the `ll` command to find which device a specific device file references:

```
# ll /dev/dsk
total 0
brw-rw-rw-   1 root    sys        7 0x0e0600 Mar 31 11:55 0s0
brw-rw-rw-   1 root    sys        7 0x0e0500 Mar 31 11:55 5s0
```

Definition

The number that starts with `0x` is the **minor number**. This number identifies the physical device that the device file references.

The general format of the Series 300/400 minor number looks like this:

`0x`*ScBaXX*

Where *Sc* is the two digit hex select code of the interface, *Ba* is the two digit device address (such as the SCSI bus address or serial port number) and *XX* is device specific information.

In this example we can see that `/dev/dsk/0s0` is a SCSI disk (select code 14 = 0x0e), bus address 6 (06). This does not match the naming convention we mentioned earlier because the primary disk of a Series 300 or 400 is always called `/dev/dsk/0s0` regardless of bus address. The disk at address 5 (`/dev/dsk/5s0`) does match the naming convention.

Reference

The *Installing Peripherals* manual for *HP9000 Series 300/400 Computers* has the complete details on select codes, device and minor number setup.

2-18. SLIDE: Series 300/400 Hardware Addressing

Instructor Notes

Purpose

Our main goal here is to introduce select codes. This information is critical when using SAM to add devices.

Key Points

- Primary disk is always named `/dev/dsk/0s0`.
- By default, only an older Series 300 would have an HP-IB disk interface at select code 14. SCSI is cheaper and faster than HP-IB.
- No two cards can have the same select code.
- The Series 400 either has 3 built-in serial interfaces (like the 425E) or it has a special 25-pin interface that is really 3 serial ports in one.

  If you have the 25-pin interface you need the `apic` driver in the kernel and the K2292 1 to 3 port adapter in order to use the other two serial ports. Their select codes are 9, 5 and 6 respectively.
- Select codes can be verified using **dmesg**.

Teaching Tips

- Show some examples of other device files.

Transition

At some point you will need to add a new device to your system. This is a simple cookbook that shows the major steps for adding a peripheral to your system.

2-19. SLIDE: Adding a Peripheral

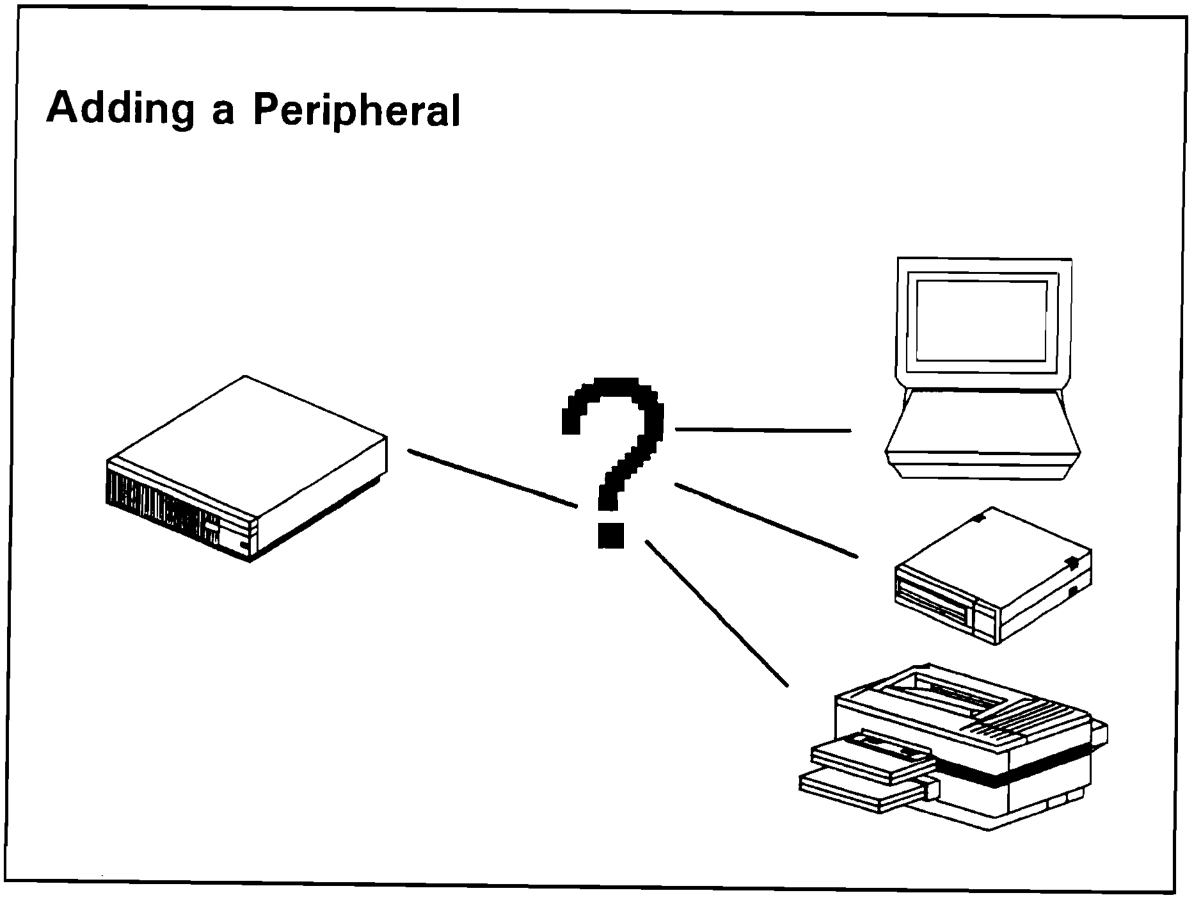

50722 2-19 26

Student Notes

The following is a generic, cookbook procedure, intended only to give an overview of the major steps involved with installing a peripheral.

On larger systems, you may choose to have your Hewlett-Packard Customer Engineer install a given peripheral, but often you can do it yourself.

Reference

To install a new peripheral, you will need the following manuals:

- The installation manual that came with the peripheral.
- The *System Administration Tasks* manual for your system.
- The *Installing Peripherals* manual for your system.

Look up the particular peripheral you're installing in *Installing Peripherals* before you start.

To add a peripheral:

1. If you are adding a card or SCSI device, shut down the computer and turn power off.
2. If applicable, set the device address on the peripheral.
3. Connect the peripheral to the computer and to the power source and turn the peripheral on.

   Consult the documentation that came with the peripheral and *Installing Peripherals* for directions.
4. If you turned off power to the computer, turn it back on and reboot it.
5. Make sure the device driver for this type of peripheral is configured into the kernel. If not you will have to reconfigure the kernel to add it.

   In general, the device driver will already be configured in your kernel and you should not have to add it.

   You can use SAM's *Kernel Reconfiguration* section to see if the driver is in your kernel and add it if it is not. Chapter 11: "Reconfiguring the HP-UX Kernel" of the *System Administration Tasks* manual explains in general how to configure device drivers into the kernel if needed. We will also discuss kernel reconfiguration later in the course.

   Installing Peripherals shows the driver name needed for each peripheral.
6. Check if a device file exists for the device. If not you will have to create the device file.

   If you use SAM to add the peripheral, SAM will make the device file for you.

   On a Series 800 the device file will be made when the system boots if it does not already exist.

   Installing Peripherals has recommendations and instructions for making device files manually if you choose not to use SAM or if SAM does not know how to make the device. For example, SAM cannot make a DDS/DAT tape drive device file on Series 3/4/700 systems. The *Installing Disk and Tape Drives* section of *Installing Peripherals* has all of the details if needed.
7. If necessary, reboot the computer.

   If you (or SAM) had to add a device driver, you will need to reboot. Otherwise, you do not need to reboot.

2-19. SLIDE: Adding a Peripheral

Instructor Notes

Purpose

This is a generic list of major steps to add a device. The details of adding some types of device can get quite lengthy so we avoid it here. Remember the audience type. They are probably on systems where all of the device are already configured. This is a reference for students that may have to add a new device. The details will be given to them in the referenced documentation.

Key Points

- A kernel reconfiguration is seldom required.
- SAM cannot add DAT drives to Series 3/4/700s. To add a DAT to a workstation the student will have to use a `mknod` command line. For example on a Series 300 or 400 this will make the device file for a SCSI DAT drive at bus address 3:

  ```
  # mknod /dev/rmt/0m c 54 0x0e0302
  ```

  On a Series 700 a DAT device file exists for a SCSI DAT at bus address 3. To make a device at address 4 the `mknod` would look something like this: SCSI

  ```
  # mknod /dev/rmt/1m c 54 0x201402
  ```

  ADD
- SAM does not add parallel printers correctly to Series 3/400 systems. The device file contains an invalid minor number (0xffff00). To properly make the device file execute:

  ```
  # mknod /dev/lp_parallel c 21 0x0c0000
  ```

- *Installing Peripherals* has all of the details for all supported devices.
- The "cookbook" procedure given in *System Administration Tasks* is wrong. It has you turning on the device before setting the device address.

2-20. LAB: Exercises

Directions

Answer the following questions on system hardware. You will have to look closely at the classroom systems, peripherals and cables to answer many of the questions.

1. List all of the SCSI devices connected to your lab system (if any). Include the bus addresses of the devices in your list.

2. List all of the HP-IB devices connected to your lab system (if any). Include bus addresses in your list.

3. What is the bus address of your system disk (the one the system booted HP-UX from)? What is its device file name?

4. List any other devices that are attached to your system. Make sure you inlcude the interface type of each device and port number if needed.

5. Draw a diagram showing the I/O interfaces and peripheral devices on the system that you administer or use at your own work site.

List all of the device files for the peripherals on your system.

6. Series 800 only:

Invoke the `ioscan -uf` command and look at the output. What are the hardware addresses of the system disk, tape drive (if any) and MUXs?

7. Series 800 only:

Run the following command and describe the output:

```
# lssf /dev/dsk/*
```

2-20. LAB: Exercises

Instructor Notes

You will have to work closely with the students during this lab to answer the specific questions about your education center equipment setup.

If you are really ambitious you could have the systems unassembled when the students walk in. Then students could get practical hands-on experience actually connecting the devices. It goes over very well with the students.

1. List all of the SCSI devices connected to your lab system (if any). Include the bus addresses of the devices in your list.

Answer:

Systems will vary but typical SCSI devices include disk drives, DDS/DAT tape drives, CD-ROM drives, and magneto-optical disk drives.

Addresses of the devices can be found on the back of the device.

Series 800 students can try the `ioscan -u` command to list devices currently connected to the system.

2. List all of the HP-IB devices connected to your lab system (if any). Include bus addresses in your list.

Answer:

Typical HP-IB devices include disk drives, DDS/DAT tape drives, cartridge tape drives, CD-ROM drives and possibly a printer.

Addresses of the devices can be found on the back of the device.

3. What is the bus address of your system disk (the one the system booted HP-UX from)? What is its device file name?

Answer:

If the system disk is SCSI its bus address should be 6.

If the system disk is HP-IB its bus address should be 0.

You can run the `bdf` command to see which device file represents the "/" directory. This can be verified by executing `ll` on the device file and looking at the minor number.

4. List any other devices that are attached to your system. Make sure you inlcude the interface type of each device and port number if needed.

Answer:

Answers will vary based on type of system and Education Center peripherals. Ask your instructor to verify your list.

5. Draw a diagram showing the I/O interfaces and peripheral devices on the system that you administer or use at your own work site.

List all of the device files for the peripherals on your system.

Answer:

Answers will vary.

6. Series 800 only:

Invoke the `ioscan -uf` command and look at the output. What are the hardware addresses of the system disk, tape drive (if any) and MUXs?

Answer:

Output of `ioscan -uf` will show the hardware addresses of all devices recognized by the system you are on (this includes adapters, modules, I/O interfaces and external peripheral devices). This command is useful if you want to get an overall picture of the hardware configuration of your system.

Ask the instructor for the specific hardware addresses for the system being used in the classroom.

7. Series 800 only:

Run the following command and describe the output:

```
# lssf /dev/dsk/*
```

Answer:

The output is the complete description of all device files in the `/dev/dsk` directory. This command can be very useful for locating device files on the Series 800.

Module 3 — System Startup and System Shutdown

Objectives

Upon completion of this module, you will be able to do the following:

- Describe the system Boot ROM startup sequence.
- Describe the difference between attended and unattended boot up and why each is used.
- List the different areas of the root disk.
- List the parts of the boot area.
- Boot your system in unattended mode (auto-boot).
- Boot your system in attended mode (manual mode).
- Explain the function of the `/etc/inittab` file.
- Identify and control the default run-level of your system.
- Describe the importance of proper system shutdown.
- Use the `/etc/shutdown` and `/etc/reboot` commands.

About This Module

This module will cover three very important aspects in the life-cycle of a computer.

System hardware boot-up

System software (HP-UX) startup and initialization

System shutdown and halt

Hardware boot-up takes different paths depending upon your system architecture and will be covered by three similar sections. Software startup and initialization and system shutdown is the same for all HP-UX systems.

Module 3 — System Startup and System Shutdown

Overview of Module 3

Our goal is to show how the system gets running, how to control some processes and how to properly shutdown the system.

We do NOT cover attended mode boot up in this course because of the time involved in revealing the details of the different architectures. If you have an audience that uses basically one type of machine, you should probably cover the attended boot up for that machine. For very diverse audiences, references are given to where they can learn more about attended mode boots.

3-1. SLIDE: What Happens During System Startup?

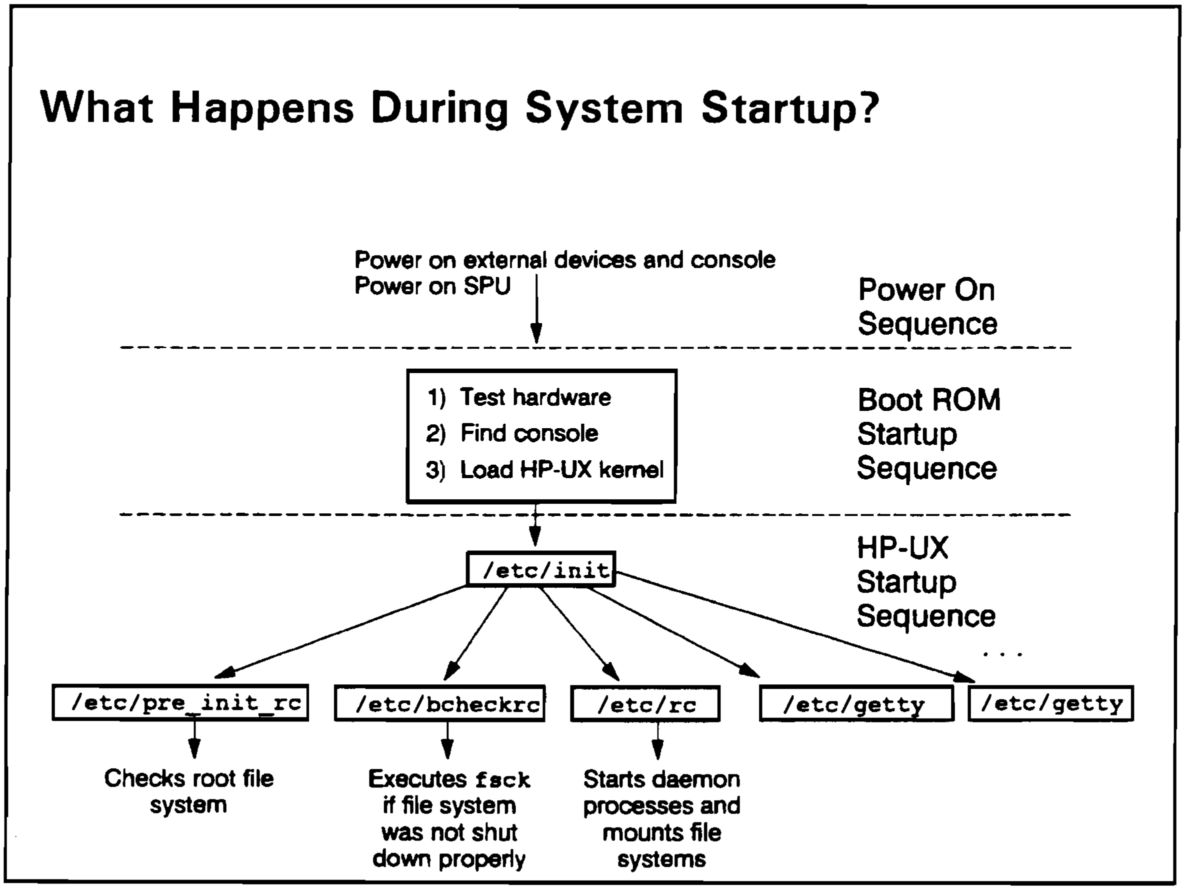

50722 3-1 27 ©1992 Hewlett-Packard Co.

Student Notes

Starting up any HP-UX system consists of three main phases:

- Power on sequence
  1. Power up all peripherals (disks, terminals, tape drives, etc.)
  2. Power on the SPU
- Boot ROM sequence
  1. Find System Console
  2. Test SPU hardware
  3. Search for bootable devices
  4. Find and load an operating system from disk or tape

Module 3 — System Startup and System Shutdown

- HP-UX startup sequence

  1. Perform system I/O and memory configuration
  2. Start the `init` process
  3. `init` runs programs to "clean up system", start system daemons and gettys
  4. Users can login and do work

In this module we will cover these three phases in more detail. Our goal is to understand enough about system startup so you can better understand what is working (and when) on your system.

3-1. SLIDE: What Happens During System Startup?

Instructor Notes

Purpose

This is an outline of the startup half of this module. We will be using this to guide the discussion of startup tasks.

Teaching Tips

- Do not go into any details of the three phases here. We will be covering them in detail on the next few slides.

Transition

Let's take a more detailed look at the three phases of system startup.

3-2. SLIDE: Power On Sequence

Power On Sequence

To properly power up your system you must:

1. Turn on all of the external disks and tape drives
2. Turn on all terminals and the system console
3. Turn on the SPU.

50722 3-2 28

Student Notes

The power-on sequence described on the slide is the proper way to power on your system. Steps 1 and 2 can be switched. The important thing is that the system console and all external peripherals are powered on and have completed their self tests before you turn on the SPU.

The console is needed to see the boot messages that are displayed during the boot ROM sequence (described next) and the HP-UX startup sequence.

The devices (especially on the SCSI and HP-IB interfaces) must be on and stable so there will be no interference when the boot ROM tests the interface. Disks and tape drives are considered stable once they have completed their power-on self tests and the front panel lights have stopped blinking.

3-2. SLIDE: Power On Sequence

Instructor Notes

Purpose

To describe the correct way to power on the system.

Key Points

- The console must be on and the external devices stable before applying power to the SPU.

Teaching Tips

- Only mention this if you have students with older Series 300s: Some older Series 300 systems with SCSI interfaces would not properly terminate the internal SCSI bus when the SPU was off. The 345 with an internal SCSI disk is one example. Certain external SCSI devices could fail their power-on self test if the SPU was off (because the bus was not properly terminated). Thus, you would have to modify the power on sequence slightly for these machines. All current machines properly terminate the bus so this is not a problem.
- Connecting all devices and the SPU to one power switch is not recommended. On some systems the disks will not complete their power-on self test before the system begins to boot. As a result, the device may not be properly recognized by the boot ROM, or worse, the system could be left in a bad state and powered off then on again.
- Many systems today have all of their disk internal to the SPU. You obviously cannot power them on first. The SPU was designed to handle this situation.

Transition

The Boot ROM takes control as soon as power is applied to the system.

3-3. SLIDE: Boot ROM Startup Sequence

Boot ROM Startup Sequence

The Boot ROM is responsible for getting the hardware up so we can boot an operating system.

The Boot ROM:

- Locates the system console
- Tests SPU hardware
- Allows attended and unattended boot modes
- Optionally configures internal SPU hardware and settings
- Searches for and Loads an operating system

50722 3-3 29

Student Notes

The **Boot ROM** (Read Only Memory) is a small machine-language program that resides off a set of computer chips in the system's SPU. When you power up the system this program is executed and causes the system to perform it's startup sequence:

1. Locate and establish a communication path with the system console device.
2. Perform internal hardware self tests and report status to the system console.
3. Search for bootable devices containing operating systems and then either allow:
   - the system to auto-boot without user interaction - **unattended mode boot** or
   - the user to choose the appropriate boot device and operating system - **attended mode boot**.

4. In attended mode the boot ROM may also allow the operator to interact with the boot ROM program to view and configure certain key values that can change the hardware/interface configuration, boot ROM configuration and the HP-UX startup sequence. **WARNING:** *This feature varies greatly with the different architectures and should only be attempted with proper supervision or guidance.*

Reference

More detailed information regarding the boot sequences of the various SPU types can be found in the manual *How HP-UX Works: Concepts for the Administrator*.

For the Series 700 and 800 you can find more information about the attended mode boot in the *HP-UX Reference Manual* under *"ISL"*, *"HPUX"* and *"BOOT_ADMIN"*.

3-3. SLIDE: Boot ROM Startup Sequence — Instructor Notes

Purpose

The goal here is to give the student an idea of the purpose of the boot ROM and to introduce attended vs. unattended boot up.

Key Points

- Over the years several different boot ROM revision changes have been made and that allow new boot features and support new boot devices. This makes it very involved to discuss the details of the boot ROM configuration modes or boot commands within the scope of this class.

Teaching Points

- Stress that the boot ROM in all models accomplishes the same tasks but the user interaction varies greatly between the various architectures.
- If the majority of your class is on one type of hardware, you may want to pull some specific information from the 5-day administration class regarding the various boot procedures. However, this is not recommended for the target audience.

Transition

The next and most complex part of the system startup procedure is the HP-UX startup sequence.

3-4. SLIDE: The HP-UX Startup Sequence

The HP-UX Startup Sequence

HP-UX:

- Initializes hardware, memory and I/O configuration
- Runs /etc/pre_init_rc
- Starts init

The init process:

- Runs system cleanup programs
- Sets the system run-level
- Runs startup programs
- Runs programs to allow users to login

50722 3-4 30

Student Notes

The very first thing the HP-UX kernel does when it starts, is to initialize the hardware, memory and I/O configuration. This includes identifying all interfaces in the SPU and possibly devices connected to the interfaces (like disks and tape drives). The kernel also locates and configures all system memory (RAM) and swap devices (disks used as part of the memory system). The kernel must also locate the disk that contains the root of the HP-UX file system. It will need the files and programs on the root file system to start everything else running. You can see this system initialization again once the system is running using the dmesg (display message) command.

Once the HP-UX kernel has completed the hardware initialization, it configures its data structures so it can run programs. The very first program to run is the /etc/pre_init_rc script.

Definition

A **script** is a command file executed using a shell (typically /bin/sh). Because these scripts are text command files, you can read them using the cat or more commands. Reading the comments in these scripts will help you to learn more about what the script does for the system.

The pre_init_rc script checks if the root file system was improperly shutdown the last time it was used. If it was, pre_init_rc will execute /etc/fsck on the root disk to resolve any potential problems. If the system *was* properly shutdown, pre_init_rc does nothing.

Next, the kernel starts the /etc/init process. The init process always has a process ID (PID) of 1 and has no parent (PPID = 0). The init process reads the /etc/inittab configuration file. This file tells init what it should do to complete the startup process. Some common things init will do at boot time include cleaning up the system from the last time it was running, checking other file system disks for potential problems, and executing the system daemons that control the print spooler, networking, etc.

On the next slide we will discuss more of what init does to control the system once the basic system startup has completed.

3-4. SLIDE: The HP-UX Startup Sequence

Instructor Notes

Purpose

Explain what HP-UX does during system startup and to introduce, `pre_init_rc` (and the concept of scripts) and init.

Teaching Tips

- If your systems are up and running you might have students look at some system scripts like `/etc/pre_init_rc` to see how they describe what they do.

Transition

We need to talk about run-levels a little to better understand how `init` works.

3-5. SLIDE: init Controls Run-Levels

init Controls Run-Levels

A run-level is a set of "things" that are happening on the system at a given time.

- At all times, HP-UX is in a particular run-level (0-6,s or S)
- All run-level definitions are in /etc/inittab
- Several run-levels are predefined (2, 3, 4, s and S)
- You can change the run-level using the init *run-level* command

```
# init 3 (Return)
```

- The who -r command will display the current run-level information:

```
# who -r (Return)
     .        run-level C  May  6 14:58    C  N  P
```

C	is the current run-level
N	is the number of times you have been in the current run-level
P	is the previous run-level

50722 3-5 31

Student Notes

At all times, the HP-UX system is in a particular run-level. A **run-level** is a system state in which a specific set of processes is allowed to run. This set of processes is defined in the /etc/inittab file for each run-level.

These run-levels are most commonly used to control user access to the system by controlling which terminals are active on the system. Typically terminals and dial-in modems are activated using the getty process (which we will discuss in the module on *Terminals and Modems*). The program that is used to control logging into the Visual User Environment (HP-VUE) vuelogin. HP-VUE (or just VUE) is a graphical windows environment using X Windows.

By cleverly designing your system run-levels you can decide when and where a set of users can access your system.

Run-levels are represented as a single character: 0-6,s and S. You can define (or change) the run-levels 0-6. You cannot change run-levels s or S because they are not part of `/etc/inittab`. Instead, they are built into the `init` command.

Your system comes with several predefined run-levels:

Run-level 2	This is the default **multi-user** run-level. This is when the system is most available to all users. Unless you change the `/etc/inittab` file, your system will be put into run-level 2 when it boots.
Run-level 3	This run-level is preconfigured for systems that have a graphics display as the system console and want to run the HP-VUE windowing environment on it. This run-level will also allow X terminals (network based graphics terminals) to run HP-VUE.
Run-level 4	This is another preconfigured run-level for systems that do **NOT** have a graphics display as their console but still want to run HP-VUE on X terminals.
Run-level s	This a special unconfigurable run-level reserved for system administration tasks. It is also referred to as **single-user mode** meaning only one user is on the system, the superuser.
Run-level S	is similar to run-level s. With `init s` only the physical system console has access to the operating system, whereas `init S` (capital S) switches the capabilities of the system console to the terminal where you are logged in, thus making it the virtual system console.

You can change the run-level of the system with the `init` command. `init` will change the run-level of the system to the level specified by its argument. `init` scans `/etc/inittab` for all entries matching the new run-level (including those entries that are valid for all run-levels) and executes the commands associated with the entries.

For example, if your system has a graphics display as its console and you want HP-VUE running there, you could tell `init` to put your system into run-level 3:

```
# init 3
```

Notice that the argument is not preceded by a - (dash) as you see with many HP-UX commands.

Whenever the run-level of the system is changed, any process which does not have an entry in `/etc/inittab` for the new run-level is sent a warning signal and then, after a 20 second grace period, is killed.

3-5. SLIDE: `init` Controls Run-Levels

Instructor Notes

Purpose

To define run-levels, how to change the system run-level and to discuss the predefined run-levels in a default `/etc/inittab` file.

Key Points

- The actions of `init` are controlled by the `/etc/inittab` file.
- Run-levels are most commonly used to control login access via the `getty` and `vuelogin` processes. `vuelogin` is started from `/etc/vuerc` which is the program we see in the `inittab` file.
- Run-levels 3 and 4 are predefined to run VUE. Run-level 3 runs VUE but no getty on the console. Therefore, we want to use run-level 3 when the console is the graphics device and will be running VUE.

  Run-level 4 runs VUE **AND** a getty on the console. This is useful if the console is **NOT** a graphics device but you have other graphics devices like a local graphics display that is not the console or X terminals. This is becoming a very popular configuration with Series 700 and 800 X terminal server systems.

  A secondary note: If you do not have a graphics display and you want VUE to run on Xterminals in run-level 4 as described, you should also edit the file `/usr/lib/X11/vue/Vuelogin/XServers`. The entry marked "local" should be removed or commented out. If you do not modify this file, VUE will try to run on the console (which is not a graphics device) and this could mess up the getty and cause other system slowdown problems. See the *HP-VUE System Administrators Guide* for more details.

Teaching Tips

- If you have workstations in the classroom with graphics consoles, you might have the students try to put their system into run-level 3 just to demonstrate how run-levels can be used.

Transition

`init` reads entries from the file `/etc/inittab`.

3-6. TEXT PAGE: The inittab File

Example Code

This is an excerpt from a typical /etc/inittab file:

```
init:2:initdefault:     — DEFAULT LEVEL
ioin::sysinit:/etc/ioinit -i >/dev/console 2>&1
brc1::bootwait:/etc/bcheckrc </dev/console >/dev/console 2>&1 # fsck, etc.
slib::bootwait:/etc/recoversl </dev/console >/dev/console 2>&1 #shared libs
brc2::bootwait:/etc/brc >/dev/console 2>&1       # boottime commands
link::wait:/bin/sh -c "rm -f /dev/syscon; \
                 ln /dev/systty /dev/syscon" >/dev/console 2>&1
cwrt::bootwait:cat /etc/copyright >/dev/syscon      # legal requirements
powf::powerwait:/etc/powerfail >/dev/console 2>&1   # power fail routines
rc  ::wait:/etc/rc </dev/console >/dev/console 2>&1 # system initialization
cons:012456:respawn:/etc/getty -h console console # Normal console mode
vue :34:respawn:/etc/vuerc                 # HP-VUE ''vuelogin'' Startup
#
a0:2:respawn:/etc/getty -h tty00 9600    # Normal Alphanumeric Terminals
a1:2:respawn:/etc/getty -h tty01 9600
     .
     .
```

Student Notes

When the system boots, it finds what run-level the system should be in using the inittab file. Whenever the system boots or the run-level is changed, the entries in the inittab file that apply to the new run-level are read and some action is taken on them. This is how we get processes to start and stop based on the run-level.

Each of the fields in the inittab file is defined as follows:

id Unique label which identifies the entry (up to four characters).

rstate Defines run-levels in which the entry will be read. An empty *rstate* field means this line is examined whenever the run-level is changed. You may have multiple run-levels in the *rstate* field which means the entry will be read when changing to several different run-levels.

action A keyword which defines how or when to execute the *process* (or program). A few of these actions are important to us:

initdefault The run-level the system will enter at boot is in the *rstate* field.

bootwait Execute the *process* only during boot up.

powerwait Execute the *process* only when returning from a power failure (Series 800 only).

off Ignore this entry.

respawn Watch the *process* that is started. If it dies, restart it.

process The command to be run if the entry's *rstate* matches the run-level and the *action* field indicates it should run.

Entries in the `/etc/inittab` file either exist when the system is first installed or will be created as needed manually by the administrator or more commonly using a tool like SAM.

Reference

Processes — Chapter 5, "Processes" *HP-UX System Administration Tasks*

Run-Levels — Chapter 6, "Run Levels" *HP-UX System Administration Tasks*

/etc/init — Chapter 6, "Run Levels" *HP-UX System Administration Tasks*
— init(1M) in the *HP-UX Reference Manual*

/etc/inittab — Chapter 6, "Run Levels" *HP-UX System Administration Tasks*
— inittab(4) in the *HP-UX Reference Manual*

3-6. TEXT PAGE: The `inittab` File — Instructor Notes

Purpose

To explain the meaning of each of the fields in `/etc/inittab`.

Key Points

- Comments are allowed in the last field and must be preceded by a '#' (the shell comment character).
- SAM will not tolerate blank lines in the inittab file. If you want to leave white space between lines use a # on the line.

Preparation Notes

You should be familiar with all of the actions in the `inittab` file, not just the ones listed in the notes. You should also have looked at all of the reference material to get a better understanding of the finer points of `init`.

Another thing to be ready for is the questions regarding all of the different programs in `inittab`. You should look up each of the processes in `inittab` and understand what they do.

Teaching Tips

- Explain to students that this is a fairly common looking inittab file.
- Only touch briefly on the `bootwait` entries.
- Use the slide to discuss the variations of the predefined run-levels.

Teaching Questions

- What would happen if the system is booted and the `inittab` file does not exist?

  (Answer: `init` will display an error and enter single-user mode.)
- What would happen to the terminals on the slide if I put the system into run-level 3 to start HP-VUE?

  (Answer: The terminals would go dead because they are not defined in run-level 3.)
- How could I make the terminals work properly in run-levels 2 and 3?

  (Answer: Add 3 to the *rstate* field for the terminal.)

Note

Terminals added with SAM are only defined to work in run-level 2. If you are using any other run-level to operate your system, you would have to modify the *rstate* entry for the terminal(s) SAM added to include the desired run level.

- How could I make this system boot up and automatically bring up HP-VUE?

(Answer: change the *rstate* field of the `initdefault` entry from 2 to 3 (or 4).

Transition

We should also talk about a few of the specific startup programs init runs during system boot up. This will also serve as a summary of how init works.

3-7. SLIDE: `init` Actions During Boot Up

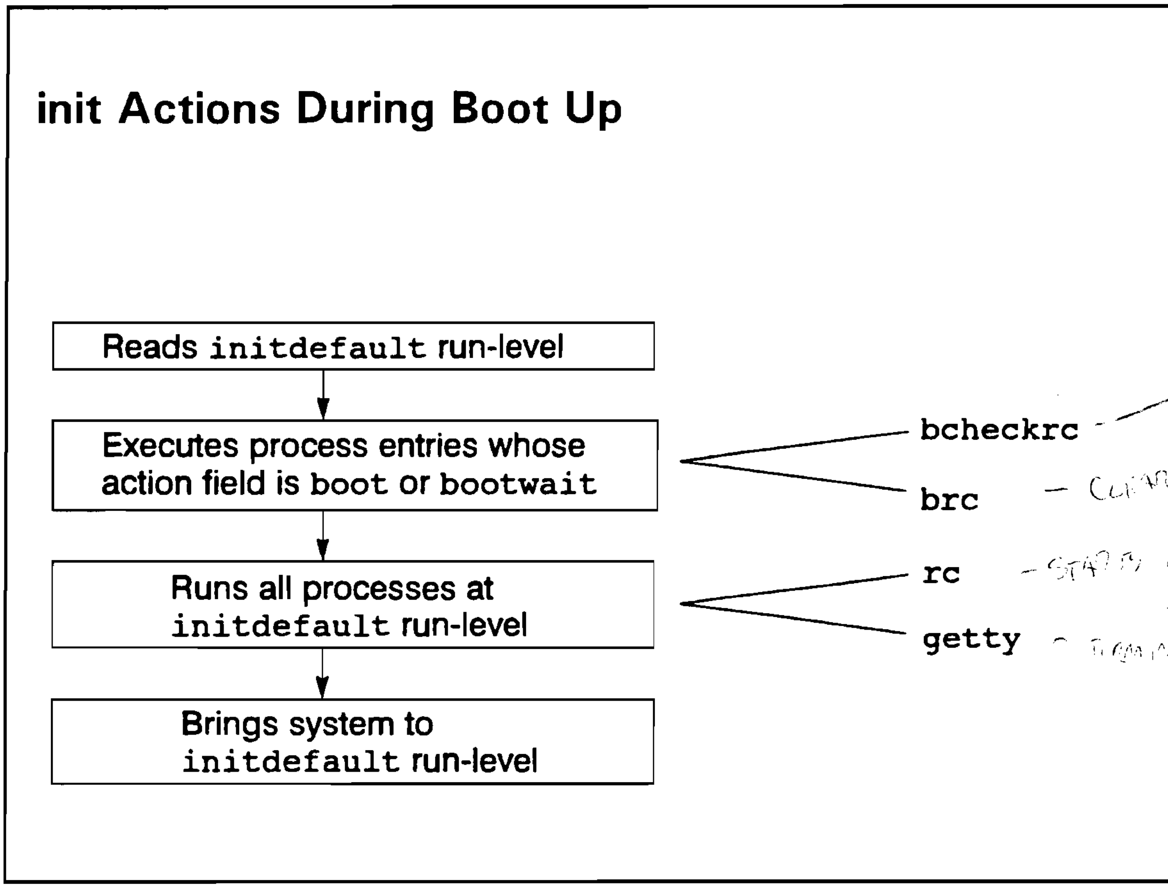

50722 3-7 32 ©1992 Hewlett-Packard Co.

Student Notes

The first thing `init` does is scan `inittab` for an `initdefault` entry in the `action` field. The *rstate* associated with this entry is the run-level `init` will enter. If there is no entry for **`initdefault`**, `init` will prompt the administrator to specify a run-level.

Next, `init` will scan `inittab` for all entries marked `boot` or `bootwait` in the *action* field. Any commands associated with these entries are executed. One of the programs that we will be concerned with is the `/etc/bcheckrc` script. This program checks the secondary file system disks to make sure they were properly shutdown. If not, `bcheckrc` will run the `fsck` tool to resolve any problems. Recall that the root file system disk was already checked before `init` started when `pre_init_rc` ran.

`init` then executes all processes associated with the `initdefault` run-level (including other entries with no *rstate* value). Typically, the system console is reset to the physical system console used by the boot ROM, the `/etc/rc` script is executed, and all login devices are activated.

3-7. SLIDE: `init` Actions During Boot Up — Instructor Notes

Purpose

To pull all of the information about `init` together into the "big picture".

Teaching Tips

- You may want students to look at the example `inittab` file or the real file on their systems to follow along with this example.
- Don't get into too much detail about the specific `inittab` entries.

Transition

We have mentioned the importance of `/etc/rc`. From the examples you should see that it is executed whenever the run-level is changed. We'll look at the purpose of this file next.

3-8. SLIDE: The /etc/rc Program

The /etc/rc Program

This line in /etc/inittab executes /etc/rc every time the system run-level is changed:

```
rc::wait:/etc/rc </dev/console>/dev/console 2>&1
```

- For all run-levels /etc/rc displays the date on screen.
- When changing from a shutdown mode into a multiuser run-level /etc/rc:
  - mounts all defined file systems – /etc /checklist
  - sets system name and time zone information
  - determines if system is standalone or in a cluster
  - prompts for current date and time
  - activates defined swap devices /etc/checklist
  - starts the printer spooler
  - starts networking
  - performs other "chores" based on system configuration

50722 3-8 33

Student Notes

As we have seen, whenever the run-level of the machine is changed with init the /etc/rc program is executed.

If the system is booting or coming back up from a shutdown state, /etc/rc takes care of starting all of the system daemons, configuring the file systems and many other startup tasks.

If you are just changing from one multiuser run-level to another (going from run-level 2 to 3, for example) /etc/rc knows that the system startup work has already been done and does not perform these tasks again.

The table below shows what functions /etc/rc performs for each system type. While not all functions are performed for each type of system, the functions are listed in the order in which they are performed. Note that some commands may not be available or installed. /etc/rc checks for the existence of all commands before attempting to run them.

Table 3-1. Functions of /etc/rc by System Type

Function	Standalone	Root Server	Diskless Cnode
set termio configuration for output device	X	X	X
create the /etc/mnttab file	X	X	X
enable quotas on root file system	X	X	X
mount all hfs volumes listed in /etc/checklist	X	X	X
preen quota statistics on hfs volumes	X	X	X
set host name	X	X	X
set the system's network name for subsequent boot ups	X	X	X
initialization: set TZ and other variables	X	X	X
set switchover host name (800 only)	X	X	X
set kind of system (standalone, server, client)	X	X	X
set the date	X	X	
save a core image of a previously crashed system	X	X	X
start swapping to all swap devices in /etc/checklist	X	X	X
start the syncer	X	X	
start the lp scheduler	X	X	
clean uucp and editor files	X	X	
start networking	X	X	X
start CSPs		X	X
start the remote boot daemon		X	X
start cron	X	X	X
start pty allocation daemon	X	X	X
start vtdaemon	X	X	X
list files found in /tmp ,/usr/tmp and /lost+found	X	X	X
clean logging files	X	X	X
start the diagnostic logging for I/O subsystem (800 only)	X	X	X
start the logging system messages (800 only)	X	X	X
start auditing processes	X	X	
Clean up old LaserROM tmp files	X	X	
local functions: any functions you put into the script	X	X	X

3-8. SLIDE: The /etc/rc Program — Instructor Notes

Key Points

- /etc/rc is one of the most important files executed by init when the system boots.
- /etc/rc decides to run the startup tasks or not based on the presence of the /etc/rcflag file. This file is created when /etc/rc first runs and tells it that the startup work has been done. The rcflag is removed when the system is shutdown or booted so when the system comes back up rc does not find rcflag and knows that the startup tasks must execute. This flag allows the administrator to add tasks to /etc/rc that need to be run every time the run-level is changed. By default, the only thing that happens every time rc runs is the date command. Read /etc/rc for more details.

3-9. TEXT PAGE: Seeing What is Running

The ps command

Use the ps (Process Status) command to see what is running on the system:

```
# <user|ps -ef| [[Return]]
     UID   PID  PPID  C    STIME TTY      TIME COMMAND
    root     0     0  0  Dec 31  ?        0:01 swapper
    root     1     0  0  Jun 23  ?        0:00 init
    root     2     0  0  Jun 23  ?        0:00 vhand
    root     3     0  0  Jun 23  ?        0:00 statdaemon
    root     7     0  0  Jun 23  ?        0:00 unhashdaemon
    root    10     0  0  Jun 23  ?        0:00 syncdaemon
    root     6     0  0  Jun 23  ?        0:00 sockregd
    root  6351     1  0 10:25:00 console  0:00 -ksh
    root    46     1  0  Jun 23  ?        0:03 syncer
    root  6357  6351  9 12:37:11 ?        0:00 ps -ef
    root   107     1  0  Jun 23  ?        0:00 /etc/cron
    root    71     1  0  Jun 23  ?        0:00 /etc/rlbdaemon
    root    76     1  0  Jun 23  ?        0:00 /etc/syslogd
    root    84     1  0  Jun 23  ?        0:00 /etc/inetd
    root    81     1  0  Jun 23  ?        0:00 /etc/portmap
    root    93     1  0  Jun 23  ?        0:00 /usr/bin/nftdaemon
    root   125   118  0  Jun 23  ?        0:00 DIAGMON
    root    99     1  0  Jun 23  ?        0:00 /etc/snmpd
    root   109     1  0  Jun 23  ?        0:00 /etc/ptydaemon
    root   118     1  0  Jun 23  ?        0:00 /usr/diag/bin/DIAGINIT
   ralph  3830     1  0 08:05:50 tty0p1   0:00 -ksh
    root   130     1  0  Jun 23  tty0p2   0:00 /etc/getty -h tty0p2 9600
    root   131     1  0  Jun 23  tty0p3   0:00 /etc/getty -h tty0p3 9600
      lp  6092     1  0  Jun 25  ?        0:00 /usr/lib/lpsched
```

Student Notes

Once the system is running you may want to see all of the active processes. The ps command is the best way to do this.

ps with no options displays the processes that you are running. This listing displays the process-id (PID), the terminal the process is running on (TTY), how much system time the process is using (TIME) and the COMMAND that started the process.

The -f option gives you a "full" listing with more information like the time the process started (STIME) and the process-id of the each process' parent process (PPID).

The -e option to ps displays everything that is running on the system.

Notice the types of things we can discover by studying the ps output on the slide:

- The printer spooler is running (/usr/lib/lpsched).
- Terminals tty0p2 and tty0p3 are both available for logging in (/etc/getty is running on them).
- The user "ralph" is logged in on tty0p1 and is running the K-shell (-ksh).
- "root" is logged in on the console and is also running the K-shell and the ps command that is producing this output (PID 6537, PPID 6531 - the PID of root's -ksh).

3-9. TEXT PAGE: Seeing What is Running — Instructor Notes

Purpose

To describe how to display the process status and infer some basic information out of the listing.

Teaching Tips

- Make sure you cover the various options of `ps`.
- Work with the students to show how we picked out the information listed in the notes from the example on the slide.

Transition

What if ralph's session was hung up and he, for some reason could not interact with his K-shell any longer? We could kill ralph's session using the `kill` command.

3-10. SLIDE: Killing Unwanted Processes

Killing Unwanted Processes

Use `kill` to remove unwanted or "hung" processes from the system:

```
# kill PID (Return)
```

50722 3-10 34

Student Notes

If you see a process that is not supposed to be running or is hung up so the user can no longer interact with it, you can remove it from the system using the `kill` command.

`kill` expects one or more process-ids identifying the processes to kill. The PIDs are found using the `ps` command.

For example, your user "ralph" comes to you and tells you his terminal is hung and he cannot do anything. You can execute `ps -ef` and look for ralph's processes:

```
# ps -ef (Return)

    UID   PID  PPID  C    STIME TTY      TIME COMMAND
    .
  ralph  1112     1 13  Jun 22  tty0p7   0:03 -ksh
  ralph 24399  1112 17 08:27:13 tty0p7   3:59 vi my_file
    .
```

Here we can see that ralph is running `vi` on `tty0p7`. You suspect that ralph messed something up in `vi` and that is the process that is hung. So, you `kill` the `vi` process.

```
# kill 24399 (Return)
# ps -ef (Return)

    UID   PID  PPID  C    STIME TTY      TIME COMMAND
    .
  ralph  1112     1 13  Jun 22  tty0p7   0:03 -ksh
    .
```

The `vi` process has been killed. If ralph had not saved his file, he can pick up almost where he left off using `vi -r my_file`. Recovering a file when the editor has been killed is a feature of `vi`.

Now, if ralph's session is still hung it must be the K-shell. So, we must kill it too:

```
# kill 1112 (Return)
# ps -ef (Return)

    UID   PID  PPID  C    STIME TTY      TIME COMMAND
    .
  ralph  1112     1 13  Jun 22  tty0p7   0:03 -ksh
    .
```

The K-shell did not go away! Many complex processes like shells will not die when killed using just the `kill` command.

If you know that a process is hung and you have tried to kill it using `kill` and it will not die; then (and only then) you can use `kill -9`:

```
# kill -9 1112 (Return)
# ps -ef (Return)

    UID   PID  PPID  C    STIME TTY      TIME COMMAND
    .
   ralph's K-shell is gone!
    .
```

Warning

`kill -9` is known as a "sure kill". You should always try killing a process with `kill` first.

If a process does not die when killed using `kill -9` the only way to remove the process from the system is to reboot HP-UX.

3-10. SLIDE: Killing Unwanted Processes

Instructor Notes

Purpose

Sometimes processes go astray. `kill` is how we remove them.

Key Points

- Always try `kill` before `kill -9`.

Transition

The last topic we need to look at to complete this module is the correct method for shutting the system down.

3-11. SLIDE: Shutting Down the System

Shutting Down the System

NEVER POWER OFF AN HP-UX SYSTEM WITHOUT PROPERLY SHUTTING DOWN. YOU COULD SERIOUSLY CORRUPT THE FILESYSTEM.

The best way to properly shutdown a running system is with the `shutdown` command:

`# shutdown [-r-h] [-y] [`*grace_period*`]`

Options:

-r	reboot the system
-h	halt the system
-y	run `shutdown` non-interactively

50722 3-11 35

Student Notes

As we have seen, most of the time the HP-UX system will be in a multi-user mode, allowing many user and system processes to run. There are occasions, however, when the administrator must change the run-level of the system from the multi-user to a single-user mode to perform sensitive administrative tasks.

For example, if the administrator wants to backup a file system, the users should not continue to work and possibly change files while the backup is occurring. The administrator should bring the system to a quiet single-user state before backing up.

Note

Executing `init s` (or S) is not the same as executing `shutdown`. `shutdown` takes extra steps `init s` does not take to make sure all disks are quiet and there are no non-essential processes running (like the networking daemons or printer spooler).

To bring the system to this quiet state simply execute `shutdown` to bring the system down in one minute. You could optionally specify a *grace_period* of seconds to allow users more or less time to log off before killing them.

shutdown will:

- Verify that the person executing `shutdown` is allowed to do so.
- Changes directory to the root (/) directory.
- Flushes all file system buffers to the disks.
- Warns users that the system is being shutdown.
- Waits *grace_period* seconds (default is 60).
- Executes all programs in the `/etc/shutdown.d` directory.
- Stops system accounting (if active).
- Kills all non-essential processes
- Detaches all file systems.
- Puts the system into a quiet single-user administrative mode.

Halting a system brings it to a complete stop. In this state, the only way to restart the system is to cycle the power or reset the hardware. You must halt the system if you are going to power it off.

Rebooting brings the system to a complete stop, but then automatically restarts the system as if you had just powered it up.

If you would like another user to have the ability to shutdown the system you can enter their name next to your system's hostname in the `/etc/shutdown.allow` file:

```
oursys root
# norton is allowed to shutdown oursys
oursys norton
```

In this example the system's name is "oursys" and the user norton is allowed to reboot or halt the system using `shutdown`. Note that norton cannot bring the system into the quiet single-user mode. He can only reboot our halt the system using `shutdown`. The system's hostname can be found by executing the `hostname` command.

Warning

If `/etc/shutdown.allow` contains any entries at all, one of them must be for `root`. If the file is empty or does not exist then only `root` can execute the `shutdown` command.

Reference

See the `shutdown` entry in Section 1 of the *HP-UX Reference Manual* for more information.

3-11. SLIDE: Shutting Down the System

Instructor Notes

Purpose

To explain what it means to "shutdown" a system. To explain why it might by necessary to shutdown a system, for example, backups, reconfiguring the kernel, and so forth. Also,

Key Points

- Turning the power off is not an appropriate method of bringing the system down. It could cause serious file system corruption.
- Make sure you explain why the `shutdown` command is used as opposed to `init s`.

Teaching Tips

- Ask the students what other activities an administrator might want to conduct with the system in single-user mode. List the responses in the student notes.

Transition

Once our system is in the quite single-user mode we can reboot or halt it quickly using the `reboot` command.

Handout: Shutting Down HP-VUE

This is a handout on the proper way to shutdown a system that is running HP-VUE. Depending on the configuration of your system, you may or may not have problems when executing `shutdown` from an HP-VUE window. The *HP-VUE System Administration Manual* recommends shutting down a system running HP-VUE using the following procedure.

Shutting Down an HP-VUE System

If your console is the graphics display the procedure to shutdown a system running HP-VUE to the quiet single-user mode is **NOT** the same as it is for other systems.

You **should not** run `shutdown` from a window in HP-VUE.

To properly shutdown a system running HP-VUE you must first change the system into a run-level where VUE is not running (like run-level 2). This stops VUE so you can then execute the `shutdown` command to bring the system to the quiet single-user mode.

This is a summary of the procedure involved:

1. Log out of the current VUE session.
2. On the login screen click the "Options" button with the left mouse button.
3. Click on "No Windows" from the pop-up options menu.
4. Login as root in text (alpha) mode.
5. Using the `init` command, change the system run-level to one that does not include VUE. Run-level 2 is a non-VUE run level by default.

```
# init 2 (Return)
```

6. Immediately exit the shell and you will see the system change into the new run-level.

```
# exit
logout root

INIT: New run level: 2

/etc/rc:

Wed Jul  8 08:56:43 EDT 1992

Console login:
```

7. Login again as root.
8. Execute the `shutdown` command as normal:

```
# shutdown
```

If you do not follow this procedure your system may get into one of several states where you may not be able to interact with the superuser shell that is started in single-user mode. This could result in an improper shutdown and potential file system corruption.

3-12. SLIDE: The reboot Command

The reboot Command

To reboot or halt the system from the quiet single-user mode use `reboot`:

`# reboot [-h] [-n] [-m` *message*`]` (Return)

No option	Reboot the system
`-h`	Halt the system so it can be powered off
`-n`	Do **not** flush file system buffers before rebooting (or halting)
`-m` *message*	Send *message* to everyone logged in before rebooting

50722 3-12 36

Student Notes

If you used shutdown to bring the system down with no options, then the system remains in the quiet single-user run-level allowing you to perform various tasks. Once you have performed these tasks, you might want to reboot or halt the system. The quickest way to do this is with the `reboot` command.

The default action of the `reboot` command is to flush the file system buffers to the disks then reboot the system. This ensures all data in memory is updated properly on the disks before the system goes down. The `-n` option **prevents the system from flushing the file system buffers to disk. Thus, data could be lost**.

Note

The only time you should use `reboot -n` is when you are specifically instructed to do so. The `fsck` tool may tell you to reboot the system using `reboot -n` after fixing the root disk.

The `-m` *message* option will display *message* at the terminals of all users logged onto the system.

Shutdown

3-12. SLIDE: The `reboot` Command

Instructor Notes

Purpose

Explain why the `reboot` command is needed and explain the syntax of the command.

Key Points

- If `fsck` modifies the root file system it will not want to take a chance of flushing potentially corrupted data from the file system buffers in memory back onto the clean file system. In this case it will tell you to reboot using `reboot -n`.

Transition

Let's look at the differences between `shutdown` and `reboot` in a little more detail.

3-13. SLIDE: shutdown and reboot Examples

shutdown and reboot Examples

- Shutdown, perform a backup, then reboot the system:

```
# shutdown

(Perform the backup commands)

# reboot
```

- Reboot with no grace period to activate a newly configured kernel (assume there are no users logged in):

```
# shutdown -r -y 0
```

- Halt the system so it can be powered off giving the users 5 minutes to log off:

```
# shutdown -h 300
```

50722 3-13 37

Student Notes

So, how do you decide whether to use shutdown or reboot? We have given you some examples on the slide. The command you use generally depends on:

- Whether users are logged in
- How quickly you need to shutdown the system

If users are logged in and doing work, use shutdown with a fairly large grace period say 5-15 minutes (300-900 seconds).

If there are no users logged in and you need to shutdown the system to do a backup, use a grace period of 0 seconds.

Use reboot if the system is already in the quite single-user shutdown mode.

3-13. SLIDE: `shutdown` and `reboot` Examples

Instructor Notes

Purpose

To point out the differences between `shutdown` and `reboot`.

Key Points

- `shutdown` is more orderly and more controlled than `reboot`. If applications are running or if users are logged in you should use `shutdown` instead of `reboot`.

Teaching Tips

- Go over the examples on the slide. Give other examples if you would like.
- Explain when users should use `shutdown` and when they should use `reboot`.

3-14. Worksession: Review Questions

1. What is the difference between executing `shutdown` and executing `init s`?

2. Why might users be angry with you if you used `reboot` to reboot the system instead of `shutdown -r`?

3. What are some of the things `/etc/rc` does at boot time?

4. How does `init` know what run-level to enter after booting?

3-14. Worksession: Review Questions

Instructor Notes

1. What is the difference between executing `shutdown` and executing `init s`?

Answer:

They both take you to single user mode but `shutdown` also takes the time to warn users, wait a specified amount of time, kill all unnecessary processes and unmount the file systems. `init s` is a single-user mode but it is not the quite single-user mode you get with `shutdown`. In fact `init s` is just one part of the `shutdown` procedure.

2. Why might users be angry with you if you used `reboot` to reboot the system instead of `shutdown -r`?

Answer:

`reboot` does not warn users of the shutdown or give them any time to properly log off. If users had applications running and work not saved, `rebooting` the system could cause them to loose their work.

`shutdown -r`, of course, will give the users a warning and a reasonable time period for the users to properly save their work before it reboots the system.

3. What are some of the things `/etc/rc` does at boot time?

Answer:

`/etc/rc`:

- mounts all defined file systems
- activates defined swap devices
- starts the printer spooler
- etc.

See *The* `/etc/rc` *Program* in this module for complete details.

4. How does `init` know what run-level to enter after booting?

Answer:

From the `initdefault` entry in `/etc/inittab`

3-15. LAB: Exercises

Directions

Your instructor will inform you as to which of the commands in this module you can try on the lab systems. Depending on the hardware available you should at least try shutting the system down and watch it boot again.

1. What is the default run-level of your system?

2. What is the current run-level of your system?

3. Determine your terminal port then find your terminal's entry in the `inittab` file.

4. If you have a graphics display as your system console, put your system in run-level 3 to start HP-VUE.

5. How can you get your system from run-level 3 back to the default run-level?

6. List all of the processes that are running on your system. Can you find your processes in the listing?

7. List just the processes you are running from your terminal. Now start a background `sleep` process that we can play with. List your processes again and make a note of the `sleep` PID.

8. Kill the sleep process then list your processes to ensure it has died.

9. Shutdown your system to the quite single-user mode using a grace period of 30 seconds.

10. Once in single-user mode list all running processes. How does this compare to the listing you saw in multi-user mode?

11. Reboot your system so it will come back up into the default run-level.

3-15. LAB: Exercises

Instructor Notes

Directions

Let your students try system shutdowns and boot ups if you have the proper equipment in your lab and the systems are available for your exclusive use (make sure another class is not being taught on the same system !).

You will, of course, need to coordinate the activities of your Series 800 users. You may have them run the shutdown simulation individually then do a system shutdown and reboot as a demonstration.

1. What is the default run-level of your system?

Answer:

Look at the `initdefault` entry of the `/etc/inittab` file.

2. What is the current run-level of your system?

Answer:

Use the `who -r` command to find your current run-level.

3. Determine your terminal port then find your terminal's entry in the `inittab` file.

Answer:

`who am i` or `tty` will allow you to see your terminal's name (device file). Scan `/etc/inittab` to find the corresponding `getty` entry.

Note If you are connected to your system via a DTC (Data Terminal Controller), you will not find an entry in `/etc/inittab`. DTC connections are handled by the `telnet` network service.

4. If you have a graphics display as your system console, put your system in run-level 3 to start HP-VUE.

Answer:

`init 3; exit` works best.

5. How can you get your system from run-level 3 back to the default run-level?

Answer:

`reboot` or `init` *default_run-level* will work.

6. List all of the processes that are running on your system. Can you find your processes in the listing?

Answer:

7. List just the processes you are running from your terminal. Now start a background `sleep` process that we can play with. List your processes again and make a note of the `sleep` PID.

Answer:

```
# ps -f
# sleep 2000 &
# ps -f
   UID   PID  PPID  C    STIME TTY      TIME COMMAND
  root 26013  9471 12 15:50:15 console  0:00 sleep 2000
  root  9471     1  6   Jun 23 console  0:05 /bin/ksh
  root 26014  9471 22 15:50:16 console  0:00 ps -f
```

The sleep PID is 26013.

8. Kill the sleep process then list your processes to ensure it has died.

Answer:

```
# kill 26013
```

9. Shutdown your system to the quite single-user mode using a grace period of 30 seconds.

Answer:

```
# shutdown 30
```

10. Once in single-user mode list all running processes. How does this compare to the listing you saw in multi-user mode?

Answer:

The listing will only have about 10 processes listed. This is why we call it the **quiet** single-user mode.

11. Reboot your system so it will come back up into the default run-level.

Answer:

```
# reboot
```

Module 4 — The System Administration Utility

Objectives

Upon completion of this module, you will be able to do the following:

- Explain the function of the System Administration Manager (SAM).
- Run SAM and navigate the menus.
- Use SAM's help facility.
- Enter information on SAM's data entry screens.
- Escape to a shell then return to SAM.

About this Module

SAM is an acronym for System Administration Manager. It is a tool that allows you to perform many system administration tasks without having to know the specific HP-UX commands that are associated with the task. SAM can also save you time and keystrokes.

Overview of Module 4

We will be using SAM a lot throughout the course. Having the students understand how SAM works now will make later sections go much more smoothly.

There is a walkthrough at the end of this chapter. Make sure you cover it either in small groups or as a class. Doing it in small groups will let the students experiment a little more but the class walkthrough will go faster.

Module 4 — The System Administration Utility

4-1. SLIDE: SAM: The System Administration Manager

The System Administration Manager

Allows you to perform many system administration tasks without having to know the specific HP-UX commands.

SAM can help you:

- Work with user accounts
- Work with peripherals (printers, modems, disks, etc.)
- Do system backups
- Reconfigure the HP-UX kernel

And much more!

50722 4-1 38

Student Notes

The System Administration Manager (SAM) is a very powerful tool designed to make your life as an administrator much easier. Most administration activities we will be discussing throughout the rest of this course can be accomplished using SAM.

Benefits

Some of the benefits of using SAM are listed below:

- 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.
- There is a robust help facility that describes the menus and data entry fields.

- You get a rich set of functions, and those functions provide significant options and control.
- You can use SAM on any HP 9000 system without relearning anything. The menus are the mostly the same on every HP-UX system. There are some differences, however, because of the architecture differences.

Strategy

We will focus on how SAM can be used to administer our systems. If a task can not be performed by SAM we will present the manual way to do it. This is a basic strategy for using SAM:

- Use SAM whenever you can to administer your system.
- Since SAM does not accommodate every task you need to perform, learn the manual procedure for performing a task.
- While performing a task with SAM, if you encounter a situation that SAM cannot accommodate, escape to a shell and perform the task manually.
- Use the manual procedure when SAM cannot perform a task or you know (as an expert) how you want to customize a functionality.

Remember, administering a system requires problem solving skills. The more you understand about your system, the better equipped you will be to solve the problems.

4-1. SLIDE: SAM: The System Administration Manager

Instructor Notes

Purpose

To introduce SAM, including the tasks it can perform and the benefits of using it.

Key Points

- Reiterate that SAM cannot be used for everything. There are some tasks SAM cannot perform, and SAM does not give you the flexibility or power that manual methods do. However, for the intended audience, SAM should perform almost all necessary administration tasks.

  When necessary, we will present the manual method for tasks SAM cannot perform. For example, SAM cannot be used to administer the disk quota functionality or add tape drives on 3/4/700 systems It cannot perform file system maintenance activities (`fsck`) either.

- The SAM utility is supplied on all HP 9000 Series HP-UX systems at version 7.0 or later. To invoke the utility, just type in sam. Also be sure that your TERM variable is set up correctly and exported.

General Notes on SAM

Note

SAM is an optionally loadable part of HP-UX. If you have not loaded SAM onto your system, you will obviously not be able to use it. You can use the update program to add SAM to your system if you did not originally load it and currently want to use it. For details about how to do this see the manual: Installing and Updating HP-UX In 8.X SAM is part of the "OS_ADMIN" update partition.

Many of the defects that had existed in previous versions of SAM have been fixed at release 8.0, making SAM a much more reliable product. The 9.0 version of SAM has both a character and OSF/Motif based interface and is said to have completely different screen layouts (so beware!).

Also, some "logging" functionality has been added to SAM. SAM users who want to investigate what actions SAM has taken can examine a logging file at `/usr/sam/log/log_file_name`. The `log_file_name` is "samlog" on Series 300s and 400s and the system hostname on Series 700s and 800s. This file was originally generated for use by SAM software developers, and is made available on an unsupported basis to users who have the desire to examine a logging file.

Warning **SAM logging is unsupported. The format and content of any information in this, or any other, log file generated by SAM is subject to change without notice.**

Here is an excerpt from SAM's log file after using SAM to add new user "mary" to an HP 9000 Series 835 system:

```
S:****** Adding user "mary".
D:         Executing (pid = 147):
            Addusr -u mary -g users -d /users/mary -i "Mary Shaw..." -s
         /bin/ksh -p "G4ZaHP2aQ.1VY" -v 201
D:       addusr: Adding user "mary" to group "users".
         password="G4ZaHP2aQ.1VY" agin g="" uid="201" gid="20" geocps=""
         shell="/bin/ksh" home_directory="/users/mary" audit_id="-1"
         audit_flag="0".
D:       addusr: Creating home directory "/users/mary" for user "mary".
D:       addusr: Calling chown(2): chown(/users/mary,201.20)
D:       addusr: Copying default files to directory "/users/mary".
D:       Child (pid = 147) terminating with status 0x0
S:****** Successfully added user "mary".
```

Transition

SAM is quite easy to use once you have the basics of how to navigate through its menus and data entry screens. The rest of this module is a short tutorial to help you get acquainted with how to use SAM.

4-2. SLIDE: Entering and Exiting SAM

Entering and Exiting SAM

```
SAM                    System Administration Manager

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

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

        File Systems ->
        Peripheral Devices ->
        Backup and Recovery ->

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

        Task Customization ...
        Other Utilities ->

        How to Use SAM

Help |  | Shell | Select Item | hpterm |  |  |  | Exit SAM
```

50722 4-2 39

Student Notes

Entering SAM

To run SAM, simply log in as root and type:

sam (Return)

The main menu is shown on the slide. Notice the different items on the menu and the softkeys (function keys) at the bottom of the screen. We will be discussing these shortly.

Exiting SAM

If you have the SAM main menu displayed, you can exit SAM by pressing the `Exit SAM` softkey (f8).

If you do not have the SAM main menu displayed, press the `Main Menu` softkey (f2) and then press the `Exit SAM` softkey (f8).

Note

If a "pop-up" window or error message box is displayed, you may have to exit the window or box before you will see the Main Menu softkey in the f2 position.

4-2. SLIDE: Entering and Exiting SAM

Instructor Notes

Purpose

To show how to invoke SAM to get the main menu and how you would exit SAM from any screen.

Note

The screen captures used in this course, both slides and figures, were taken from an HPTERM terminal using Windows. You may need to point out to the students that their screens might have a slightly different configuration of softkey buttons.

4-3. SLIDE: Using SAM's Menus

Using SAM's Menus

Several keys help you get around in SAM:

- Tab, ▼, ▶ — Move to next item
- Shift Tab, ▲, ◀ — Move to previous item
- Return or `Select Item` — Select currently highlighted item
- `Previous Menu` — Go up one level in the menu hierarchy
- `Main Menu` — Return to SAM's main menu
- `Help` — Display a help screen on the highlighted item

50722 4-3 40

Student Notes

When you enter SAM, you will notice that the first menu item `Users ->` is highlighted. When you press the `Select Item` softkey f4, the item which is currently highlighted will be the item which is "selected" from the menu. Pressing Return on a a highlighted menu item will also select it.

Use the arrow keys or the Tab key on your keyboard to change which item is highlighted. The down arrow (▼), the right arrow (▶), and the Tab key will highlight the next item in the menu. The up arrow (▲), the left arrow (◀), and the Shift Tab key will highlight the previous item in the menu.

An arrow "—>" at the end of a menu item means you will get another menu when the item is selected. When a menu item ends in " ... " it means that the item will display a **data entry screen** where SAM will prompt you for more information.

When you enter one of SAM's sub-menus you will notice that f8 displays `Previous Menu`. This is so you can traverse backwards through the menus if needed.

Notation Conventions

For the remainder of this course, when we discuss traversing menus and screens in SAM we will use the following conventions:

Notation	Meaning
`Menu Item`	A menu item to select on one of SAM's menus.
Key	A key to be pressed.
Key key 2	Two keys to be pressed simultaneously.
`Softkey Name`	A softkey on the bottom of the SAM screen.
"*Screen Name*"	The name of one of SAM's screens.
"Field Name"	The name of a field to be completed on one of SAM's data entry screens.
`Data item`	Value of a data field.
<u>user supplied value</u>	Value you should enter into a data field.

4-3. SLIDE: Using SAM's Menus

Instructor Notes

Purpose

To describe how to traverse SAM's menus.

Key Points

- The actions in SAM are fairly intuitive. Most of the time, you traverse SAM with keys that make sense.

Transition

Let's take a look at one of SAM's data entry screens.

4-4. SLIDE: Entering Data in SAM

Entering Data in SAM

```
SAM                     Add a New User Account to the System

   Fill in or modify the desired fields and then press "Perform Task".

   Login name . . . . . . . . . . . _____

   Primary group name . . . . . . . users______

   Home directory . . . . . . . . . /users/_____________________

   Start-up program . . . . . . . . /bin/sh_____________________

   Real name . . . . . . . . . . . ___________________________(optional)

   Office location . . . . . . . . ________________ (optional)

   Office phone  . . . . . . . . . ________________ (optional)

   Home phone  . . . . . . . . . . ________________ (optional)

   Modify user's defaults (y or n) n

 Help | Main | Shell | Perform | hpterm |  |  |  | Exit
      | Menu |       | Task    |        |  |  |  | Task
```

50722 4-4 41

Student Notes

The slide shows one **data entry** screen we will see in SAM. This screen is used when we want to add a new user account to our system. We will not discuss the details of adding a user here. We just want to discuss the basic data entry format in SAM.

The Tab, Shift-Tab, and arrow keys (left arrow, right arrow, up arrow, down arrow) function in a manner similar to how they work on menu screens, with the following exceptions:

- Within a data entry field, the left and right arrow keys move the cursor within the field, except when they are at the beginning or end of a field, where they will attempt to move you to the previous or next field.
- On data entry screens, the (Return) key functions in the same way as the (Tab) key.

- When you attempt to leave a data entry field, validation checks are often performed to verify that the data entered in the field is of the correct form. If the value in the field is not correct, a window will appear describing what is wrong and what you should do to correct it.
- You can use the (Back Space) and editing keys ((Insert Line), (Delete Line), (Insert Char) and (Delete Char)) to edit information in any field.

Notice that some of the fields on the data entry screen are already filled in. These are default values SAM assumes we want to use. You can type over the default values.

Also note that some of the fields are marked as (optional). SAM does not require any data in this field. These fields are usually used as additional information for the administrator.

If you wish to leave the form without actually doing anything you can use the `Exit Task` softkey (f8). This will return you to the screen that originally produced the form.

If SAM has only one or two small pieces of information, it will not give you a full screen to fill in. Instead it will simply display a window over the current screen that prompts you for the information. You can either enter the data SAM is requesting or press `Exit Window` softkey (f8) to cancel the operation.

4-4. SLIDE: Entering Data in SAM — Instructor Notes

Purpose

To show the different types of fields on a data entry screen and how to traverse and edit them.

Key Points

- Note the key traversals as listed in the student notes. Especially note that (f8) will usually get you out of what you are doing.

Transition

What if you don't know what kind of value SAM is looking for on a menu or data entry field?

4-5. SLIDE: Getting Help in SAM

Getting Help in SAM

```
SAM                    Add a New User Account to the System
    Fill in or modify the desired fields and then press "Perform Task".

                              Help: Login Name
    DEFINITION:    This is the name your system uses to identify the user. The login
                   name is also referred to as the "user name". Each time the user
                   logs in to your system, (s)he must supply this name.

                   An individual user may have more than one account and
                   corresponding login name.

    HOW TO GET:    Normally, the user's first name, last name, or a combination of
                   first and last name is used. The login name must have the
                   following characteristics:
                        * Must begin with an alphabetic character.
                        * Can include up to 8 alphanumeric characters.
                        * Cannot contain blank spaces.
                        * Cannot already exist on the system.

                   — Press "Exit Help" to return to the form. —

                                   hpterm                                  Exit
                                                                           Help
```

50722 4-5 42

Student Notes

If you are ever unsure of what value you need to put in a data entry field or what a menu item will do, position the cursor in that field or on the item and press the `Help` softkey (f1).

A window will be displayed describing what needs to be entered in the field. In many cases, SAM will present a list of valid options for you to choose from as part of the help screen.

One of the values will be highlighted. You can use the arrow keys or the (Tab) key to change which item is highlighted (just as you would with menu screens). If there are more entries than will fit in the display at one time, you will see up arrows (^^^^^^^^) or down arrows (vvvvvvvv) indicating that there are entries above or below the ones that are currently displayed. There are many places in SAM where you will see these arrows. As you continue to move the highlighted item up or down, the entries in the display will automatically scroll. When you have highlighted the value you want to select, press the `Select Item` softkey or (Return). The value you selected will be placed in the field you wanted help on.

If you did not want to select any of the displayed items, you could press the `Exit Help` softkey (f8). The help window would disappear and the value in the data entry field would remain unchanged from the value that was present when you pressed the `Help` softkey.

On SAM's main menu you will also see a menu choice called `How to Use SAM`. This is a brief description of how to use the basic SAM features.

4-5. SLIDE: Getting Help in SAM — Instructor Notes

Purpose

SAM has an excellent context sensitive help facility. It is important that students know how to use it for both understanding what they are doing and for displaying valid options where available.

Key Points

- SAM help is context sensitive. It will display help on any specific item.
- If there are a finite number of valid choices SAM will list them and you can select one.
- The `Help` softkey can be used as a compliment the other HP-UX manuals.

Transition

Once we have a data entry screen filled in we would like SAM to actually perform the action.

4-6. SLIDE: Having SAM Perform the Task

Having SAM Perform the Task

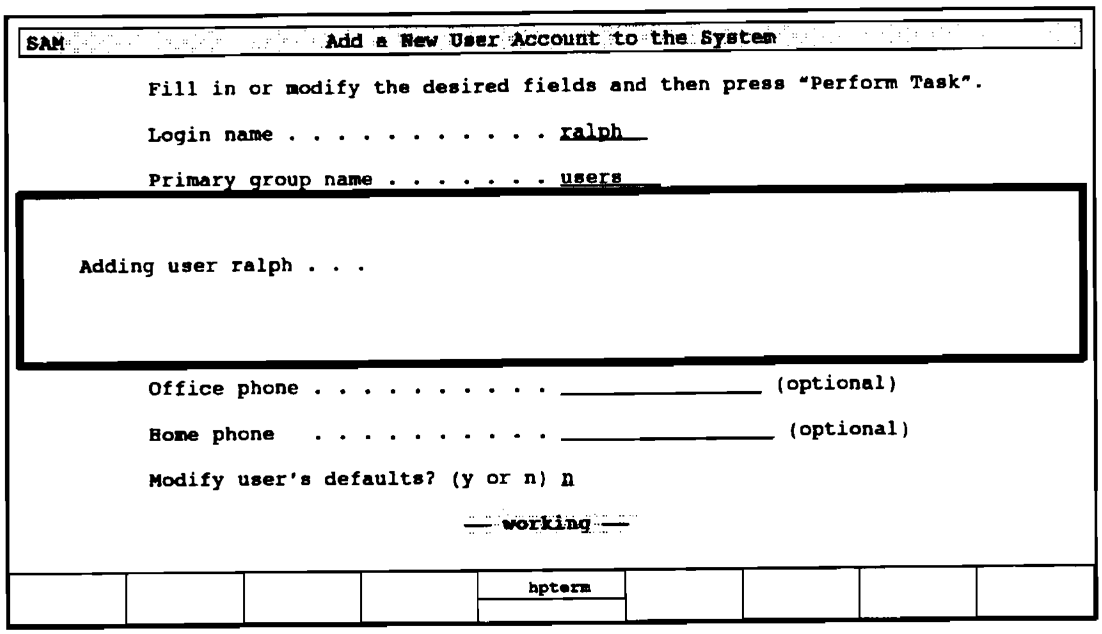

50722 4-6 43

Student Notes

Once you have completed the data entry screen, you can have SAM perform the task for you by pressing the `Perform Task` softkey (f4).

You will see status messages in SAM windows as the task is being performed. While SAM is working on the task, you will also see the "`-- working --`" flag at the bottom of the screen.

If SAM encounters an error while performing a task, that error and any relevant information will also be displayed in a window so you can take corrective action.

When SAM has completed the task you will see a "Task Completed" window. At this point you are either returned to a blank data entry screen so you can do the task again (so you could add a second user, for example) or you may be returned to the menu.

```
SAM              Add a New User Account to the System

     Fill in or modify the desired fields and then press "Perform Task".

     Login name . . . . . . . . . . . . . ralph

     Primary group name . . . . . . . . . users

Task completed.

ralph has been added to the system.

                         — Press the space bar to continue. —

     Office phone . . . . . . . . . . . . ______________ (optional)

     Home phone . . . . . . . . . . . . . ______________ (optional)

     Modify user's defaults? (y or n)  n

                                              hpterm
```

Add a New User Account

4-6. SLIDE: Having SAM Perform the Task — Instructor Notes

Purpose

To describe how to make SAM act on a task and what a user should expect while SAM is working on the specified task.

Key Points

- Status windows should be read carefully. An error condition may slip by if you simply respond to a window without knowing what it was for. This is a surprisingly common problem among SAM users.

Transition

There is an easy way to break out of SAM temporarily to run HP-UX shell commands without loosing your place in SAM.

4-7. SLIDE: Escaping to a Shell

Escaping to a Shell

```
Type "exit" to return to SAM.

# ll /users
total 14
drwxr-xr-x   2 ralph    users      1024 Jul  8 16:32 ralph
drwxr-xr-x   2 norton   users      1024 Jul  8 16:30 norton
# ll -a /users/ralph
total 12
drwxr-xr-x   2 ralph    users      1024 Jul  8 16:32 .
dr-xr-xr-x   9 bin      bin        1024 Jul  8 16:32 ..
-rw-r--r--   1 ralph    users       818 Jul  8 16:32 .cshrc
-rw-r--r--   1 ralph    users       347 Jul  8 16:32 .exrc
-rw-r--r--   1 ralph    users       377 Jul  8 16:32 .login
-rw-r--r--   1 ralph    users       382 Jul  8 16:32 .profile
#
#
#
```

50722 4-7 44

Student Notes

Occasionally, you might find it necessary to momentarily exit SAM to run an HP-UX command, perhaps to gather some information for a data entry field in the middle of one of SAM's data entry screens. Rather than backing all the way out of SAM to get the information and then once again entering SAM and traversing the menu structure to the point you left off, you can press the `Shell` softkey (f3) to start an HP-UX shell.

SAM's user interface is temporarily suspended while the shell is executing. When you are finished with the shell, type exit and you will be returned to SAM's user interface (on the screen where you were, prior to running the shell).

4-7. SLIDE: Escaping to a Shell

Instructor Notes

Key Points

- When you exit the spawned shell you return to SAM exactly where you left off.

Transition

Some tasks in SAM can be tailored to better fit your needs.

4-8. SLIDE: Customizing SAM

Customizing SAM

```
SAM                         Task Customization

      SAM can perform additional processing before and/or after the tasks listed below
      by executing programs you specify. This allows you to customize these tasks to
      meet your requirements. Press "Help" on any item to learn how SAM performs the
      task or to view the parameters that will be passed to your executable. Enter the
      name of your executable(s) and press "Perform Task".

                  SAM Task                         Program to Run (full path)   When Run

      Add a new user account to the system:        _________________________    before
                                                   _________________________    after
      Remove a user account from the system:       _________________________    before
                                                   _________________________    after
      Add cluster clients:                         _________________________    before
                                                   _________________________    after
      Remove cluster clients:                      _________________________    before
                                                   _________________________    after

 Help | Main | Shell | Perform | hpterm |   |   |   | Exit
      | Menu |       | Task    |        |   |   |   | Task
```

50722 4-8 45

Student Notes

There are currently four tasks in SAM that you can customize:

- Add a new user account to the system
- Remove a user account from the system
- Add cluster client
- Remove cluster clients

SAM will let you specify a program you wrote to run before and/or after each of the listed tasks. This can be useful if you need to copy key application files into a new user's home directory. To do this simply create a small shell program with the appropriate commands in it then put that program name on the task customization screen to execute after a new user has been added. Now SAM will execute your program after it adds any user to the system.

To see an example of how to construct the program SAM will execute before or after adding a new user, see the file `/usr/sam/config/ct_adduser.ex`.

The customizations for adding or removing cluster clients deal with systems configured in an HP-UX Cluster. If you will be working with HP-UX clusters, this may interest you. The sample program for cluster node task customization is `/usr/sam/config/ct_addnode.ex`.

4-8. SLIDE: Customizing SAM

Instructor Notes

Purpose

Many administrators of application systems will like the ability to customize how users are added. Especially the ability to copy configuration files into a new users directory automatically.

The ability to customize how cluster nodes are added/removed will not typically be as interesting to them.

Key Points

- You will need some shell programming experience to write the task customization program. If the user is not comfortable with (or doesn't know anything about) shell programming but still wants this customization, they can contact HP for some consulting time.

Transition

Now that we have seen what SAM can do and how to use SAM let's try a quick SAM session.

4-9. LAB: SAM Walkthrough

Directions

This is a short tutorial to help you get acquainted with how to use SAM. We will be using SAM to add a new user account to the system. We will not be focusing on user administration here, instead we will be concentrating on the traversal of SAM and using its many features.

Only one person at a time (per system) should use SAM. If more than one person or lab group is logged into the same machine you must coordinate your activities somehow.

For series 800 students: if the pseudo-shell is available, each pseudo superuser group can do the labs simultaneously. Ask your instructor if and how the pseudo super user accounts are set up on the series 800 used for class.

Entering SAM

To start up the System Administration Manager, use the command:

```
# sam (Return)
```

This will display SAM's main menu on your screen.

Selecting a menu item

1. Try changing the highlighted item using the arrow keys and (Tab) key now.
2. When you are comfortable with changing which item is highlighted, move the highlight back to the menu item `Users ->`.
3. Press the `Select Item` softkey (f4).

You have now selected the *"Users"* menu. Recall that when a menu item ends with the character sequence "`->`" it means that selecting this item will display another menu.

4. From the *"Users"* menu, be sure that the first item `Add a New User Account to the System...` is highlighted and press the `Select Item` softkey (f4). Notice that the menu choice ends with " ... ".

Entering Data

You should now see the data entry screen that allows you to add a new user account. When a menu item ends with the character sequence " ... " it means that the item will display a data entry screen.

5. Enter the value "%%%" in the "Login name" field and press (Return).

   A window will be displayed indicating that an invalid character has been entered.
6. Press the space bar to make the pop-up window disappear.

To complete the data entry screen, do the following:

7. In the "Login name" data entry field, type frunelda.

   Because the name "frunelda" fills the complete data entry field, the cursor will automatically move to the next field (after validating that "frunelda" is a valid user name and that there isn't another "frunelda" on your system already).

   If the value you enter doesn't fill an entire data entry field, press (Return) to move to the next field.

   When the cursor moves to the second data entry field ("Primary group name"), notice that the value in the "Home directory" data entry field changes to be /users/frunelda.
8. Press (Return) twice to accept the default values in the "Primary group name" and "Home directory" fields.
9. In the "Start-up program" data entry field, use the right arrow key to position the cursor over the letter "s" in the default value /bin/sh.
10. Press ksh on your keyboard. The value in this field should now be /bin/ksh.

    If it is /bin/kshsh, your keyboard was in insert-character mode. Simply press the (Delete char) key twice to remove the second sh.
11. For this example, leave the optional data entry fields blank.

Getting Help in SAM

If you are ever unsure of what value you need to put in a data entry field, position the cursor in that field and press the `Help` softkey (f1).

12. Use the arrow keys to move the cursor back to the "Start-up program" data entry field (if it isn't already there).
13. Press the `Help` softkey.
14. While the "*Help: Start-up Program*" help screen is displayed, use the up and down arrow keys to highlight the various choices of startup programs.

    Don't do this, but if you did not want to select any of the displayed items, you could press the `Exit Help` softkey (f8). The help window would disappear and the value in the data entry field would remain unchanged from the value that was present when you pressed the `Help` softkey.
15. Highlight the value "/bin/csh" and press (Return).

    Notice that the value in the data entry field changes to /bin/csh.

Having SAM perform the Task

Once you have completed the data entry screen, you can have SAM perform the task for you by pressing the `Perform Task` softkey (f4).

Do this now, for the user you are adding:

16. Press the `Perform Task` softkey.

    SAM will then ask you to assign a password for this user.

17. Press the `Help` softkey to display information about valid passwords. Notice that entering the sequence `,..` (comma, period, period) will allow the user to assign their own password to their account the first time that they log in.

18. Press the `Exit Help` softkey to exit the help screen.

19. Enter `,..` in the password field and press Return.

    Notice that as you enter the password, it is not displayed. This is to prevent someone who is looking over your shoulder from seeing the password that you assign.

    Because the password is not displayed when you type it in, SAM will ask you to re-enter the password (to confirm it).

20. Enter `,..` in the field again (to confirm the password).

    SAM will now add the user to your system for you. As SAM performs the task you will see status message windows on the screen. When it is finished, it will let you know and ask you to press the space bar to continue.

    The fields on the data entry screen will be cleared. If you have further users to add you could complete the data entry screen again with the data for the next user. In this case, we are finished.

21. Press the `Exit Task` softkey f8 to return the "*Users*" menu.

22. Now, see if you can remove the user "frunelda" from your system.
    - When you are asked if you want to remove all of the users files from the system, answer y (yes).
    - When you have removed the user and SAM requests the name of the next user to remove, press the `Exit Window` softkey f8 to return to the "*Users*" menu.

Escaping to a Shell

You can suspend SAM from any point where the f3 softkey is labeled `Shell`.

1. Press the `Shell` softkey f3.

   An HP-UX shell will be started. The same one you get when you log in as root.

2. At the shell prompt, type:

   ```
   # ps (Return)
   ```

   The ps command lets you view the processes that you have running in your login session. Notice the process `samx`—it is the suspended SAM process.

3. We can also check to see that the user "frunelda" has really been removed from the system by typing this command at the shell prompt:

   `# ls /users/frunelda` (Return)

   You should get an error because the directory no longer exists.

4. When you're finished with this shell and ready to return to SAM, type:

   `# exit` (Return)

   The shell will terminate and SAM will resume.

Exiting SAM

1. If you do not have the SAM main menu displayed, press the **Main Menu** softkey (f2).

2. Once you have the SAM main menu displayed, you can exit SAM by pressing the **Exit SAM** softkey (f8).

4-9. LAB: SAM Walkthrough

Instructor Notes

Purpose

We will be using SAM a great deal throughout the class. It is important to get the students comfortable with traversing SAM, entering data and getting help.

Teaching Tips

- This can either be done as a class walk through or as small teams (1-3 per system). The latter gives the students more opportunity to play around and see what else SAM can do.
- If you are using a Series 800 you may want the students to use the pseudo superuser environment (P-shell) so they can all run SAM simultaneously. If you only have a few 800 users (1-3) it is easiest to just let them log in as root and use SAM as a team.

Transition

Now that we know how to use SAM, let's put it to work!

Module 5 — User Accounts

Objectives

Upon completion of this module, you will be able to do the following:

- Use SAM to add users.
- Use SAM to remove users and their files safely.
- Use SAM to deactivate and reactivate user accounts temporarily.
- Communicate with users on the system.
- Explain the login process.
- Set up a user's login shell environment.
- Explain some basic user account security concerns.

About This Module

In this module we will look at how the System Administration Manager (SAM) can be used for adding and removing users. We will also discuss the many ways an administrator has of communicating with the users, basic user security considerations and setting up a user's login environment so their applications will behave properly (a very important and often ignored aspect of user administration).

Overview of Module 5

This is very important information to all operators or administrators. This module covers all important information regarding user administration.

Groups and group administration are not covered in this module. It is assumed that once a user is assigned to a group (when the user is added with SAM) they will not be changing their group identity. Most installations dedicated to running a single application only use one group for everyone. If you feel that groups are important to the students, you should cover a little on groups at the end of the chapter. A quick example of the use of groups is given in the instructor notes.

5-1. SLIDE: User Administration with SAM

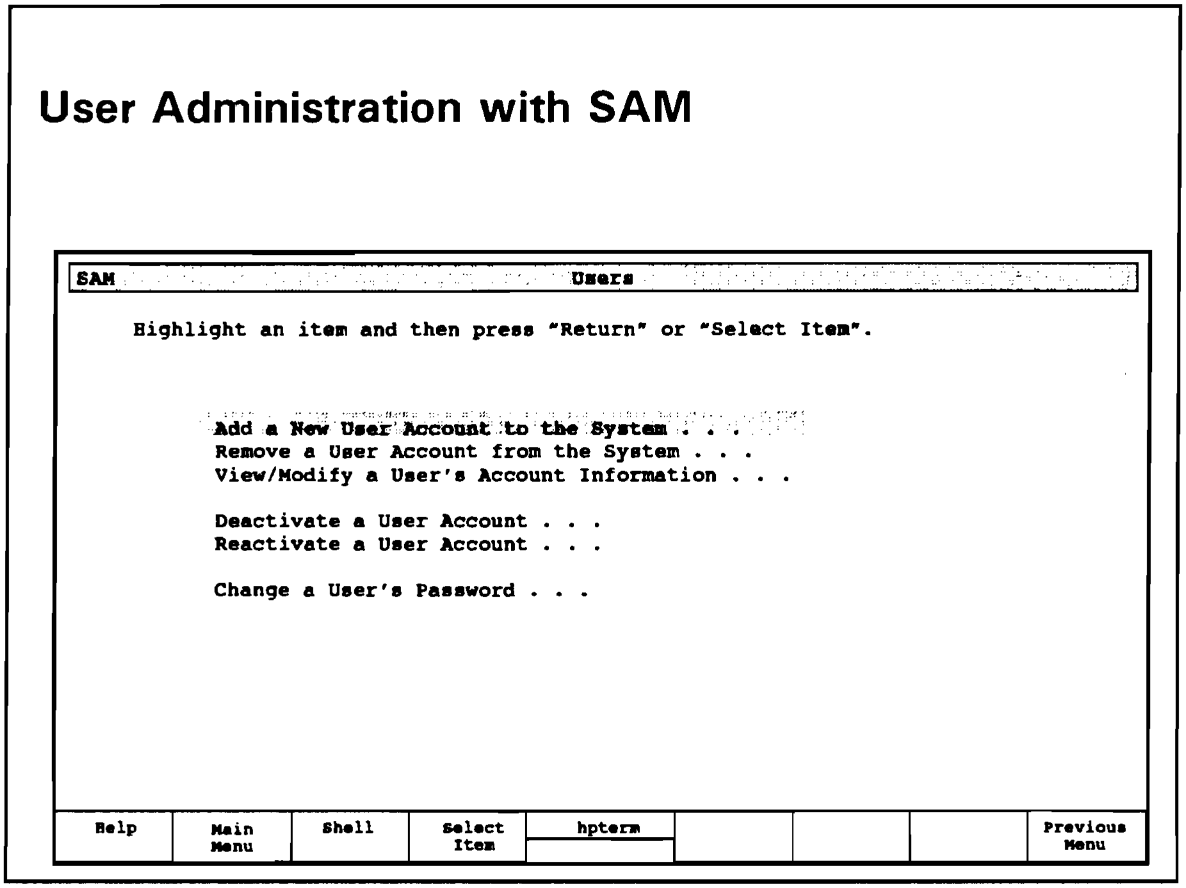

50722 5-1 46

Student Notes:

You will get the menu on the slide when you select `Users ->` from the main SAM screen. This screen shows the many user-related tasks that SAM can perform.

- `Add a New User Account to the System ...` makes all changes to system configuration files then builds the user's login directory and login files. This does everything needed to allow a user access to the system.
- `Remove a User Account from the System ...` is just the opposite. You have a choice to remove the user's files and directories or not.
- `View/Modify a User's Account Information ...` allows you to change things like the user's login program or password.

- If you `Deactivate a User Account ...` , you deny that person access to the system until you `Reactivate` it.
- `Change a User's Password ...` does just that.

Note

The ... after each menu item indicates that you will get a data entry screen to fill in the appropriate information. The actual task will not be performed until the information is given and you tell SAM to perform the task.

5-1. SLIDE: User Administration with SAM Instructor Notes

Purpose

This slide also shows the students the user-related tasks SAM can perform. This sets us up for the discussion on the next few slides.

Teaching Tips

- Don't go into detail on any one of the tasks SAM can perform here. This is meant to be a broad overview of SAM's user administration abilities.

5-2. SLIDE: Adding a User with SAM

Adding a User with SAM

```
SAM                  Add a New User Account to the System

    Fill in or modify the desired fields and then press "Perform Task".

    Login name . . . . . . . . . . .  norton
    Primary group name . . . . . . .  users
    Home directory  . . . . . . . .   /users/norton
    Start-up program  . . . . . . .   /bin/ksh
    Real name  . . . . . . . . . .    Norton                 (optional)
    Office location  . . . . . . .    Late night TV      (optional)
    Office phone . . . . . . . . .                        (optional)
    Home phone . . . . . . . . . .                         (optional)
    Modify user's defaults? (y or n)   n

 Help | Main Menu | Shell | Perform Task | hpterm |  |  |  | Exit Task
```

50722 5-2 47

Student Notes

To add a new user follow these steps:

```
# sam (Return)

    Users ->

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

If you enter an incorrect value in any of the data entry fields SAM will display an error window and allow (make) you to correct the error.

Each of the first four pieces of information is required for SAM to properly add the new user:

- "Login Name" - This is the name your system uses to identify the user. The login name is also referred to as the "user name". Each time the user logs in to your system, the user must supply this name.

  Usually, the user's first name, last name, or a combination of first and last names is used. The login name must have the following characteristics:

  - Must begin with an alphabetic character.
  - Can include up to 8 alphanumeric characters.
  - Cannot contain blank spaces.
  - Cannot already exist on the system.

- "Primary group name" - Users are organized into groups to facilitate users accessing other group members' files. For a more complete discussion of the relationship between groups and file permissions, see the manual "How HP-UX Works: Concepts for the System Administrator".

  If you are not sure what groups are configured on your system, the list of groups can be displayed using the `Help` softkey. You can either select one of these groups or enter a new group name. If you enter a new group name SAM will configure that group into the system.

- "Home directory" - When a user logs into the system, (s)he is placed in the specified directory. The home directory usually has the same name as the user's login name, and is typically located under the `/users` directory.

  SAM creates the directory you specify and copies several login files into it that will set up the user's working environment when (s)he logs in.

- "Start-up program" - When a user logs in to the system, a "shell" is started. This shell (program) is the interface between the user and HP-UX.

  You can select the default shell (`/bin/sh`) or use the `Help` screen to see what other shells are supplied with the system.

The four optional fields allow you to add extra information about the users, such as the user's full name, telephone extension, organization, or building number. This can help you locate users if you need to communicate with them. These fields are also used by several commands that need the real users information.

You should not have to "Modify user's defaults?" so just accept the default answer (n). If you answer (y) to this question you can use the `Help` screens to learn more about the other user attributes you can change.

Once you have filled in the necessary fields, press the `Perform Task` softkey to begin adding the user.

Reference

All of the information you see on the "Add a New User Account" data entry screen is stored in a file called `/etc/passwd`. This is the file that is referenced whenever the system needs information regarding a user account. To learn more about `/etc/passwd` look at the `passwd` entry in section 4 of the *HP-UX Reference Manual*.

Information regarding user groups is stored in the file `/etc/group`. To learn more about this file look at the `group` entry in section 4 of the *HP-UX Reference Manual*.

5-2. SLIDE: Adding a User with SAM — Instructor Notes

Purpose

The intent here is to show the process flow of adding a new user with SAM. We also want to introduce the important information needed to administer user accounts.

Key Points

You may get questions about the user defaults screen depending on the curiosity level of the students. Stress that none of these things need to be changed to successfully add a user. Other key points to stress about the user defaults screen are listed here for your reference.

- Change User's uid Number. The user identify (uid) number is used by the file system to identify users. Owners of files in a file system are determined by the uid numbers, not by the login names of the users.

  The user identify (uid) number is automatically assigned when the user is added to the system. You might want to manually give a user a specific uid for the following reasons:

  - Give the user "root" permission (uid=0).
  - Make the user's uid match that user's uid on other networked file systems.

  If the uid matches an existing uid, then the user will be warned of the conflict and asked to confirm the intention to use the uid anyway.

- Log in with X11 Windows. A new user can be set up so that the X11 Windowing system is automatically started when that user logs in to the system. Then, when the user exists the window system, she or he is logged out automatically.

  Before checking "y" on this item, ensure that X11 windows is installed on your system (check that `/usr/bin/x11start` exists). If not, use `/etc/update` to add it to your system (it is in the `X11-RUN` fileset in the `WINDOWS` partition).

Note

This feature is not necessary and will not work correctly if the HP-VUE or xdm X Window login environments are being used.

This will also not work correctly if the console is NOT a graphics display (which will be the case for almost all Series 800 machines).

- Log in with Terminal Session Manager. Terminal Session Manager (TSM) is a multi-session manager (windows) for ASCII terminals. A new user can be set up so that TSM is automatically started when that user logs in to the system. Then, when the user exits TSM, she or he is logged out automatically.

Depending on the software configuration purchased, TSM may be an optional product. Make sure that TSM is installed on your system (check that `/usr/bin/tsm` exists). If TSM is not installed you can load it using `update` (load the "TERM-MNGR-MIN" fileset from the "OS_FEATURES" partition).

A Note on Groups

If students do not understand the use of groups, you might show them a quick example:

If we have a system that both research and marketing teams use we may decide to protect the marketing information from the research team and vice versa using groups. In this example we could have two groups `mktg` and `resrch`. This alone does not do much for security. At this point we would simply see that any files or directories created by a member of the `resrch` group would be owned by that group (using an `ll` command). The same for the `mktg` group.

To actually protect the files and directories owned by one group from the other group we will have to change the permissions also. For example, if there were a directory called `/project` that we wanted to be private to the members of the `resrch` group we could set ownership and permissions as follows:

```
# chown root /project
# chgrp resrch /project
# chmod 770 /project
# ll -d /project
drwxrwx---   4 root     resrch       2048 Jul  8 07:31 /project
```

Now only members of the group `resrch` can access files and directories in `/project`. Members of other groups, including `mktg` will receive `permission denied` errors if they try to access the directory.

If one member of the `mktg` group needed access to the `/project` directory, they could be made a member of both groups and switch their group identity using the `newgrp`(1) command.

Transition

The last piece of information SAM needs to add our user is a password.

5-3. SLIDE: Setting a User Password

Setting a User Password

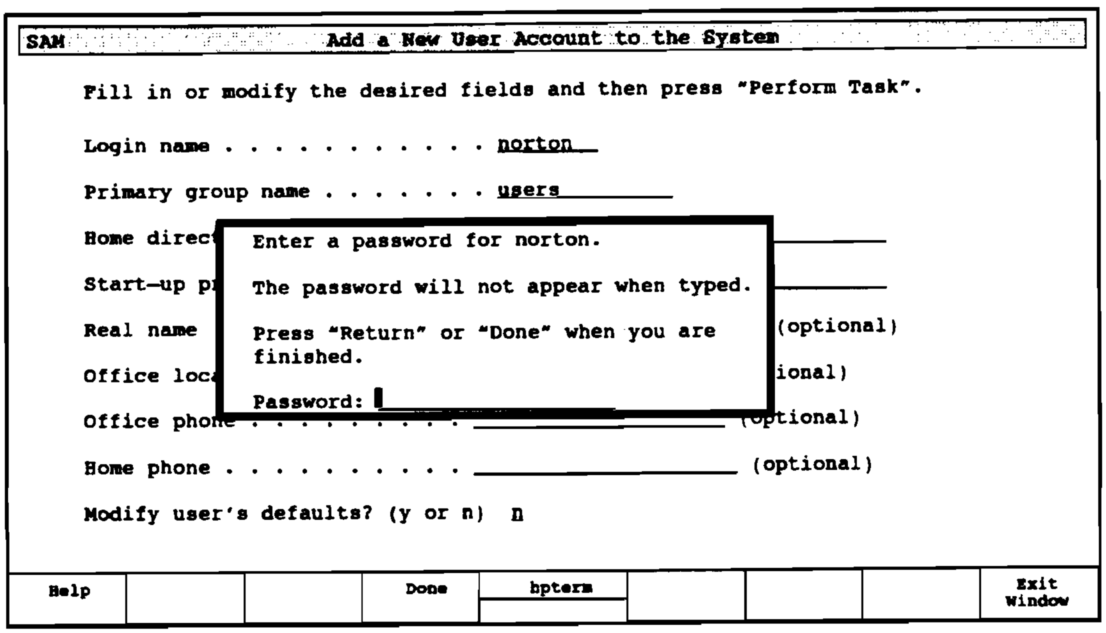

50722 5-3 48

Student Notes

When you finish filling out the data entry screen and press `Perform Task` you will be prompted for a password to be placed on the new user's account. The password you type will not appear on the screen so SAM will ask you to type it a second time to verify the password.

You have several choices for a user password:

- Any string of characters or numbers — The user will have to enter the exact string when (s)he logs in.
- The three characters ", . ." (comma period period) — This will force the user to set a password the first time they log in. The password they enter will become their permanent password for logging in.

- (Return) — This leaves the password blank. The user will not be prompted for a password when they log in.
- * — This asterisk deactivates the user account. The user can not log in until a valid password is assigned (changed to one of the above choices).

Changing the User's Password

If you want to change the user's password at a later time you can select `Change a User's Password ...` from the "*Users*" menu. SAM prompts you for the user name to change then the new password as described above.

You can also change any user's password using the **passwd** command. For example, the command

```
# passwd ralph
```

allows you to change the password for user "ralph".

The password you assign can be changed by the user using the **passwd** command. A number of rules apply when a **user** changes a password. The password a user sets must:

- contain at least 6 characters and only the first eight characters are significant.
- contain at least 2 alphabetic characters (uppercase or lowercase).
- contain at least 1 numeric (0-9) or special (-, _, $, etc.) character. Avoid using characters such as (Back space), #, @ in the password because they may have special meaning at login.
- not be any variation of the login name.
- differ from the old one by at least 3 characters.

Note

These rules on password characteristics do not apply to the superuser when changing a user's password.

Reference

See the **passwd** command in section 1 of the *HP-UX Reference Manual*.

The SAM help screens are also a good source of help (as always).

5-3. SLIDE: Setting a User Password

Instructor Notes

Purpose

To describe how SAM can set and administer user passwords and to mention the other ways the password can be changed using the passwd(1) command.

Key Points

- The password can be set or changed in these ways:
  1. In SAM when adding a user.
  2. In SAM to explicitly change a users password.
  3. The superuser executing the `passwd user_name` command.
  4. The user executing the `passwd` command.

Transition

Let's look at some of the other things SAM can do with user accounts.

5-4. SLIDE: Removing User Accounts

Removing User Accounts

```
SAM                    Remove a User Account from the System

    Press "Perform Task" to remove this user's account. Do you want to remove all
    files and directories belonging to this user from the system? (y or n)  n

          Login name . . . . . . . . . . . norton       Account status: ACTIVE

          Primary group name . . . . . . . users

          Home directory  . . . . . . . .  /users/norton

          Start-up program  . . . . . . .  /bin/ksh

          Real name  . . . . . . . . . . . Norton

          Office location  . . . . . . . . Late night TV

          Office phone . . . . . . . . . .

          Home phone . . . . . . . . . . .

 Help | Main Menu | Shell | Perform Task | hpterm |  |  |  | Exit Task
```

50722 5-4 49

Student Notes

If someone leaves your company or department you may want to take their user account off the system to prevent anyone from using the account.

To remove a user account from the system follow these steps:

```
# sam (Return)

    Users ->

        Remove a User Account from the System ...
```

SAM will prompt you for the user to be removed and display the screen you see on the slide.

You must decide if you want to remove the user's files or not. If you choose to remove the user's files SAM will start a process that searches the disk for all files owned by the specified user and will remove them.

If you choose to keep the user's files SAM will bring up a window and ask what you want to do with the user files:

```
Mark one of the choices below with an "x" and press "Done".

x  Assign all of the user's files and directories to another user.
_  Remove only the user's home directory and files within it.
```

Mark your choice with an x then press the `Done` softkey.

SAM will verify your choice and prompt you to continue. When you decide to continue, SAM will complete the task by removing the user from the `/etc/passwd` and `/etc/group` files and taking the specified action on the user's files. SAM will start a background process to remove the files (if specified) and to make sure the deleted user no longer appears in the access control list (ACL) of any file.

```
Task completed.

norton has been removed from the system. A process has been activated in
the background to remove all ACL entries for this user.


                  — Press the space bar to continue. —
```

Once you press the space bar, SAM will then prompt you for the next user to remove. If you do not want to remove another user account simply press the `Exit Window` softkey to return to the "*Users*" screen.

Note

SAM will not let you remove critical users such as root or bin from the system. It will also prevent you from removing critical files in directories such as /bin or /etc.

Use the `Help` softkey when you first enter the "*Remove a User Account*" screen to learn more on how you can protect other users and files from accidental removal.

5-4. SLIDE: Removing User Accounts — Instructor Notes

Purpose

To describe some of the details involved with removing users and their files from the system.

Key Points

A mechanism exists in SAM to protect system administrators from inadvertently removing users and files belonging to removed users. You may choose to discuss this more completely with the students.

When removing users or files from a system, there is always the unfortunate likelihood that the wrong user may be removed or that files belonging to a user who is removed are deleted inadvertently during the removal process.

For example, the user `bin` is the owner of the majority of the executable commands on the system. Removing this user would obviously be disastrous. On the other hand, suppose user joe owns all of the files comprising the test suite for a project. It may be appropriate to remove joe, but the test suite should be left intact and assigned to a new owner.

SAM provides two features to help protect against inadvertent removal of users or files when removing users:

- Editable list of users to exclude from removal.

  When prompting for the name of a user to remove from the system, SAM checks the name given against a list of names specified in the file `/usr/sam/config/rmuser.excl`. If the name matches one within the file, SAM does not remove the user.

- Editable list of files to exclude from removal when a user is removed from the system.

  When SAM removes a user, all files (or a subset thereof) for that user are also removed, unless the ownership is given to another user. Before removing a file belonging to the user, SAM checks to see if the file resides in a path that has been excluded from removal. SAM uses the file `/usr/sam/config/rmfiles.excl` to determine which paths have been excluded from removal. So, for example, if the path `/users/joe/test` is named in the file, SAM will not remove any files residing beneath that directory. SAM logs a list of all files it removes in the file `/tmp/sam_remove.log`.

You can edit the files `/usr/sam/config/rmuser.excl` and `/usr/sam/config/rmfiles.excl` to contain users and directories that you want to exclude from removal from SAM.

Here is a sample /usr/sam/config/rmuser.excl file:

```
root
daemon
bin
adm
lp
hpdb
```

Here is a sample /usr/sam/config/rmfiles.excl file:

```
/bin
/dev
/etc
/lib
/usr
/SYSBCKUP
/SYSDEBUG
/hp-ux
```

Transition

If you just want to keep a user from logging in without removing their account, you can simply deactivate the user account temporarily.

5-5. SLIDE: Deactivating and Reactivating a User

Deactivating and Reactivating a User

```
SAM                     Deactive a User Account

    Press "Perform Task" to deactive this user's account. Do you want to perform any
    special processing on this user's files? (y or n)  n

          Login name . . . . . . . . . . . ralph

          Group name . . . . . . . . . . . users

          Home directory . . . . . . . . . /users/ralph

          Start-up program . . . . . . . . /bin/ksh

          Real name  . . . . . . . . . . .

          Office location  . . . . . . . .

          Office phone . . . . . . . . . .

          Home phone . . . . . . . . . . .

 Help | Main | Shell | Perform | hpterm | | | | Exit
      | Menu |       | Task    |        | | | | Task
```

50722 5-5 50

Student Notes

Deactivating a user account will keep the account on the system but the password is changed on the account so the user can not log in. This is useful if someone is on leave or vacation and you want to protect their account from anyone else using it.

```
# sam (Return)

    Users ->

        Deactivate a User Account...
```

This will bring up a small window where SAM asks you for the user to be deactivated. When you supply that name you will see the screen on the slide with all of the user's information filled in.

If you want to keep the user's files and directories in tact just press Perform Task and the user account will be deactivated.

SAM, however, will also let you do special processing on files owned by the specified user. The options for file processing are the same as for removing a user. You can:

- Assign all of the user's files and directories to another user (SAM will prompt you for the new owner's name).
- Remove only the user's home directory and files within it.
- Remove all files and directories belonging to this user.

Once again, a background job will be launched to complete the file processing and SAM will prompt for the next user to be deactivated.

To **Reactivate a User Account ...** (when the user returns from vacation, for example) you must provide SAM with the user's name and then assign a password as described earlier.

5-5. SLIDE: Deactivating and Reactivating a User

Instructor Notes

Purpose

This is an underutilized feature of SAM. Relatively few people do this but the option is there on the menu and we should describe a little on what it does.

Key Points

- You can take action on the files of the user to be deactivated. You will NOT usually want to do any special processing on the files because we presume the user is returning someday.

Teaching Question

What is the difference between removing a user and deactivating a user?

Transition

Now that we have all of these users, how can we keep in touch with them?

5-6. SLIDE: Communicating with System Users

Communicating with System Users

The `news` Command

- Displays files in `/usr/news` directory
- Used for long, less than critical messages

The `mail`, `mailx`, and `elm` Commands

- Mailer programs used to send messages to specific users
- Users read mail by invoking the specific mailer

The `wall` Command

- Sends a message to all users logged in
- Immediately interrupts whatever the user is doing

50722 5-6 51

Student Notes

You will frequently need to communicate with the users of the system. The importance of the message and who needs to be aware of it should be considered when selecting the communication method to use.

The `news` Command

For messages that are not of great importance ("nice to know" as opposed to "need to know"), create a file containing the news item and place it in the `/usr/news` directory. The name of the file should be the subject of the news item. When a user logs in, if there is a new entry in `/usr/news`, the user sees the following message:

`news:` *news_filename*

The user can then read the message with the `news` command.

The news facility can be used to effectively communicate things like backup schedules, printer configurations or locations, news items relevant to the whole department, etc.

Mailers

If you need to send a message to an individual user, then use one of the mailers. We will not cover the usage of the mailers in this class because this is probably a review for most students. If you are unfamiliar with using mailers you should try `elm`. It is a menu driven mailer that works with the `vi` editor. `elm` is probably the easiest mailer to use. Other mailers include `mail` and `mailx`.

The `wall` Command

To simultaneously send a message to all users logged in, use the `wall` (write all) command. This command is typically used to generate a message that is of immediate concern to the users. For example, if, for some reason, the system must be shutdown immediately, `wall` can be used to send a message to all the users currently logged in warning them of the impending shutdown.

`wall` is typically run interactively where you type in the message to be sent from the command line.

```
# wall
The system will be shutdown in 5 minutes.  Please log off.
Ctrl d

Broadcast Message from root (console) Sat Mar 18 11:22:43...
The system will be shutdown in 5 minutes.  Please log off.
```

For more complex messages you may decide to create a file containing the message the run `wall` telling it to get the message from the file.

```
# wall < message_file
```

`wall` can only be invoked by the superuser.

5-6. SLIDE: Communicating with System Users Instructor Notes

Purpose

To explain to the students the different options they have for communicating with users. *Continued on next slide!*

Key Points

- Users will not be informed of new news items if they are using HP-VUE because the windows do not typically run the normal login files. The "console" window is one exception to this rule because it is a "login shell" window.
- `/etc/motd` and `/etc/issue` will not work for the same reason. To perform a similar task in HP-VUE see the file `/usr/lib/X11/vue/Vuelogin/Xsession` and the `vuehello` command.

Transition

Let's continue on the next slide with `/etc/issue` and "/etc/motd".

5-6. SLIDE: Communicating with System Users (Continued)

Communicating with System Users (Continued)

The etc/issue File

- Contents are displayed just before `login:` prompt
- Used for short welcome message

The etc/motd File

- Contents are displayed when a user logs in
- Used for short important messages
- Typically changed on a daily basis

50722 5-6 52

Student Notes

The /etc/issue File

The contents of this file will be displayed on the screen just before the `login:` prompt. This file usually contains a short welcome message and system identifier.

The /etc/motd File

For messages that every user should be aware of, place an entry in /etc/motd (message of the day). For example, if you decide to shut the system down for an evening for preventive maintenance work, place a message in /etc/motd letting the users know when the system will be down so they can plan their work time accordingly.

The contents of the /etc/motd file will be displayed every time a user logs in.

5-6. SLIDE: Communicating with System Users (Continued)

Instructor Notes

See previous slide for KEY POINTS, etc.

Transition

Next, we want to cover some of the basic security concerns regarding user accounts.

5-7. SLIDE: User Security Considerations

User Security Considerations

- Set proper permissions on files and directories
- Protect passwords
- Log off unwanted terminals
  - Set time out (`TMOUT`) variable if available
  - Use `lock` to lock an unattended terminal
- Monitor security logging files
  - `/etc/wtmp` — logs successful login attempts
  - `/etc/btmp` — logs unsuccessful login attempts
  - `/usr/adm/sulog` — logs use of `su` command
  - `/etc/securetty` — specifies the `tty` files on which the root can log in

50722 5-7 53

Student Notes

As an administrator of an HP-UX system, maintaining the overall security of the system is a primary responsibility. There are many aspects of system security, some of which we will touch upon here.

File and Directory Permissions

One of the biggest security concerns is using proper file and directory permissions. This is a quick review of how these permissions work.

The permissions are set and changed using the `chmod` command:

`chmod` *mode path_name*

Where *mode* is a usually a set of three digits each between 0 and 7, and *path_name* is the name of the file or directory to be changed.

The first digit represents what the user (owner) of the file/directory can do to it. The second represents what the members of the group the file belongs to can do to it. The third digit represents what anyone else can do to it.

Each digit is actually a sum of the permission values:

Permission	**Value**
Read (r)	4
Write (w)	2
Execute (x)	1

Simply add the values of the permissions you want set. For example, read and write permissions would be 4 + 2 = 6, just read permission would be 4. To make a file readable and writable for the owner but only readable by anyone else, you would use *mode* 644.

```
# ll myfile
-rw-rw-rw-   1 ralph    users      53788 Jun 18 10:11 myfile
# chmod 644 myfile
# ll myfile
-rw-r--r--   1 ralph    users      53788 Jun 18 10:11 myfile
```

The User's Responsibility

The individual user should assume some responsibility for security. Users should:

- Set proper permissions on their files.
- Protect their passwords.
- Log out when they leave their terminal.

Each individual user must be responsible for maintaining the correct permissions on their files and directories. Correct permissions are the only way to prevent users from accessing, changing or removing other users' material. Users and the administrator should periodically check the permissions set on various important files and directories using the `ll` command.

Each individual user must also be responsible for protecting his or her password. Users should be encouraged not to divulge their password to anyone, and should clearly understand the implications of supplying their password to another person.

Users should also be responsible for logging out of the system while their terminals are unattended. One way you, as an administrator, can encourage this is to use a time-out value such as the `TMOUT` variable in the Korn Shell (or the `autologout` variable in the C shell). `TMOUT` can be set to the number of seconds a Korn shell (ksh) can remain idle (no keyboard activity at a shell prompt) before the shell dies and the session goes away. You can set default values of `TMOUT` and `autologout` in the system login files (which we will discuss shortly).

Alternatively, a user can use the lock command to lock his or her terminal. When invoked, lock requests a key. The key must be entered twice for verification purposes.

```
$ lock
Key:
Again:
LOCKED
```

The terminal is locked until the key is entered again. This allows a users to leave temporarily without logging off.

If using the HP-VUE X Windows environment, a user can lock all of their windows at once by clicking the left mouse button on the picture of the padlock at the bottom of the screen. The user will have to enter their login password or the superuser password to unlock the display.

Tips for the Administrator

You should make sure that permissions on system files and directories are correct using the ll command. The following list shows you what the permissions should be on some specific files and directories:

/etc/	755 (rwxr-xr-x), owned by root.
/bin/*	555 (r-xr-xr-x), owned by bin.
/etc/passwd	444 (r—r—r—), owned by root.
/users/user_name	755 (rwxr-xr-x), owned by user_name.

where user_name represents a user's login name and /users/user_name is that user's login directory. Permissions of 755 on a directory prevent anyone but the owner of the directory from creating files in or removing files from that directory.

Note

X

The rm command checks directory permissions NOT file permissions. If a user has write permission on a directory they can remove any file in that directory regardless of the file permissions!

The only exception to this rule is that a normal user can not remove a read-only file owned by root, even if the directory is writable.

Some Security Log Files

Several security logging files exist on the system. They will help you find if anyone is trying to break system security. Here are just a few of these files:

`/etc/wtmp` — The system uses this file to keep a history of logins, logouts, system state and date changes. You can view a list of user login sessions using the `/etc/last` command.

`/etc/btmp` — If this file exists, the system uses it to keep track of **bad** login attempts. Use the `/etc/lastb` command to access the contents of this file.

`/usr/adm/sulog` — If this text file exists, the system uses it to keep track of information on the use of the `su` command. It contains the name of the user attempting to switch their id using `su` and user name they are switching to. The terminal name, the date and time, and whether the attempt was successful To view the information, use the `more` or `cat` command.

Note

`/etc/wtmp` is usually created for you by the system. `/etc/btmp` must be created manually. If either of these files do not exist they can be created using the touch command:

```
# touch /etc/btmp        (or /etc/wtmp)
# chmod 600 /etc/btmp
```

Note

All three of these files will continue to grow as time goes on. They could become very large if not emptied periodically. This is a simple method for emptying the log file while keeping a recent copy:

```
# cp /etc/wtmp /etc/OLDwtmp
# > /etc/wtmp
```

These commands could be run once a month to trim the log files (this is most easily done using `cron` which will be discussed later).

5-7. SLIDE: User Security Considerations — Instructor Notes

Purpose

To make the students aware that security for the system is their responsibility. They should encourage their users to follow the practices on this page.

Key Points

- Permissions are sometimes confusing to users at first. The permissions review gives them a chance to see it once more. Make sure you cover it.
- Good system security is a collaborative effort between system users and the administrator. Students should be highly encouraged to log out from (or `lock`) their sessions when they leave their terminal.
- Correct file permissions are very important. System files should generally be owned by root with no write access for anyone else.
- `rm` looks at directory permissions NOT file permissions!
- Protecting home directories with 755 mode is critical for network security because of the .rhosts file (see hosts.equiv(4)).
- `/etc/wtmp` will only be created automatically if the accounting software is installed on the system. If not, wtmp will have to be created manually as discussed.

Transition

The last item in user administration we want to cover is customization of the user environment.

Handout: Introduction to Access Control Lists

The following is a brief introduction to access control lists (ACLs). This is a flexible extension to the normal "user/group/other" permissions on UNIX systems. It is quite common for students to ask how they can allow/deny access to a particular file or directory. The normal permissions can not allow a user-by-user control, ACLs can. The next page can be photocopied for the students.

Introduction to Access Control Lists (ACLs)

ACL's allow you to extend the normal "user, group, other" file permissions to allow or restrict access to files and directories to any number of individual users or groups. This type of control is impossible with the standard UNIX permissions and the chmod command.

chacl(1) is the command used for setting/removing access. The general form of chacl is:

```
chacl user.group=mode file [file ...]
```

The lsacl(1) command lists access control lists on files:

```
lsacl file [file ...]
```

lsacl also lists the normal UNIX permissions in ACL format:

```
$ ll myfile
-rw-rw-rw-   1 ralph    users         99 Mar 18 18:37 myfile
$ lsacl myfile
(ralph.%,rw-)(%.users,rw-)(%.%,rw-) myfile
```

The following command allows the user norton from any group (% means any) read and write access to myfile:

```
$ chacl norton.%=rw myfile
```

A long listing only shows that there are now extensions to the normal permissions by displaying a "+" at the end of the permissions list:

```
$ ll myfile
-rw-r--r--+  1 ralph    users        524 Mar 18 18:38 myfile
```

lsacl, however, displays the full access control list:

```
$ lsacl myfile
(ralph.%,rw-)(norton.%,rw-)(%.users,r--)(%.%,r--) myfile
```

To remove all permissions from myfile (nothing for the mode) for any user in the group students:

```
$ chacl %.students=  myfile
$ lsacl myfile
(ralph.%,rw-)(norton.%,rw-)(%.users,r--)(%.students,---)(%.%,r--) myfile
```

The "-d" option will delete an access control. In this example we are taking out the controls on norton from myfile:

```
$ chacl -d norton.% myfile
```

```
$ lsacl myfile
(ralph.%,rw-)(%.users,r--)(%.students,---)(%.%,r--) myfile
```

Now norton has no special access controls and will abide by what is listed for his group or the "other" permissions (%.%)

NOTE: If you have an ACL on a file then you use chmod on that file, all ACLs will be lost.

5-8. SLIDE: The User Environment

The User Environment

The current environment can be displayed using the env command:

```
# env
LANG=C
PATH=/usr/bin/X11:/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin
MANPATH=/usr/man:/usr/contrib/man:/usr/local/man
WINDOWID=37748753
EDITOR=vi
LOGNAME=ralph
MAIL=/usr/mail/ralph
DISPLAY=viper:0.0
SHELL=/bin/ksh
HOME=/users/ralph
TERM=hpterm
PWD=/users/ralph/project
TZ=EST5EDT
COLUMNS=80
LINES=24
```

50722 5-8 54

Student Notes

Recall that when a user logs in a startup program is executed allowing them to interact with the system. This startup program is typically one of the UNIX shells. Associated with each shell you run is your **user environment** (sometimes call the **shell environment**).

To see your current environment you can use the env command as shown on the slide. As you can see, your environment describes many things about your working environment to the system and the programs that you run.

An analogy to your user environment is your office environment. In the office, characteristics such as lighting, noise and temperature are the same for all workers. the factors in your office that are unique to you make up your specific environment. These factors include your desk location and size, your mail box location, your name plaque, etc. Your specific office environment is unique to you just like every user on the system has a shell user environment that is unique to them.

The typical shell user environment contains many things. Each piece of information is identified by a name:

Variable	**Value**
HOME	The name of your home directory.
MAIL	Your electronic mail file.
TZ	The time zone you are working in.
PATH	The list of directories where your shell will search for the commands you enter.
TERM, *COLUMNS*, *LINES*	Your terminal type and size.
LOGNAME	The name you used to log in.
DISPLAY, *ENV*, *EDITOR*, etc.	Many other things the applications you run may need.

A simple example of how your shell environment is used is with commands like `vi` and `more`. Whenever you run commands like these, they look in your environment to find what type and size terminal you are using so they can format the output correctly.

Action

As a system administrator it is your job to set up the default user environment so your users can function and interact properly with their applications.

If you are using an application it is important to read the application installation notes to see what (if any) environment variables need to be configured to properly run the application.

5-8. SLIDE: The User Environment

Instructor Notes

Purpose

Students should already know something about the user environment from the Basics I or Fundamentals prerequisite classes. This is intended as a review for students that may not have seen the environment before or took the other classes some time ago.

Key Points

- Very few applications run correctly without a properly configured environment.
- The application installation notes will usually describe the necessary environment variable settings.
- Almost all applications (CAD, databases, etc.) require some modification to the user environment.

Teaching Tips

- Stress the importance of the environment as a motivation for the next few slides.

Transition

There are several ways to customize the users' environment.

5-9. SLIDE: Customizing the User Environment

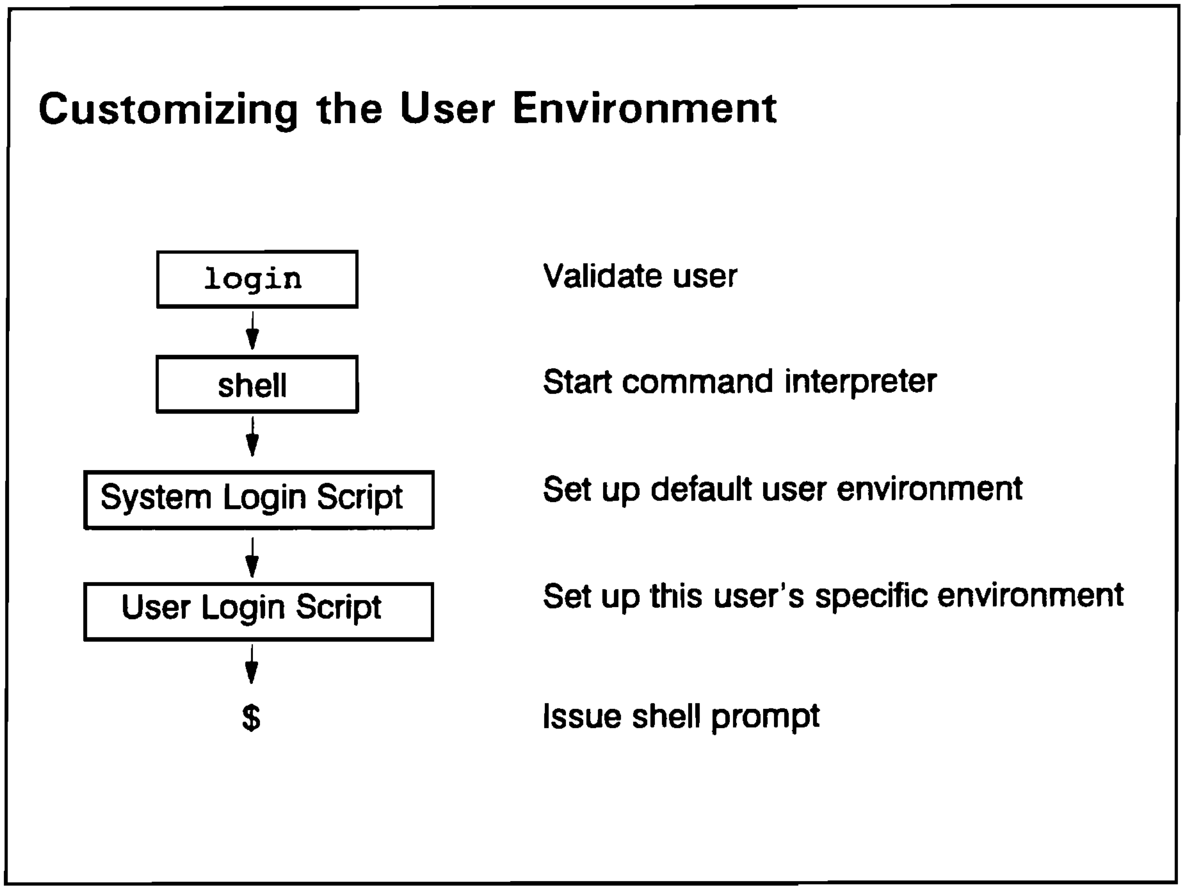

50722 5-9 55

Student Notes

To set up its environment a shell uses a series of **login scripts**. There is typically one script that the administrator can edit that will affect the environment of all users on the system and there is a second script that can be controlled and customized by the end user to suit their unique needs if necessary. The shell prompt is issued after all login scripts have run.

System Login Scripts

The startup shell runs the appropriate **system login script**, which initializes the default user environment:

If the Shell is ...	The System Login Script is ...
Bourne (/bin/sh)	/etc/profile
Korn (/bin/ksh)	/etc/profile
Restricted (/bin/rsh, /bin/rksh)	/etc/profile
C (/bin/csh)	/etc/csh.login

As shipped, these scripts define and export the default values of the environment variables PATH, TZ, and TERM.

Since these scripts are run for all users at login, the system administrator can modify these files to set up the default environment for all users.

User Login Scripts

After the system login script is run, the shell runs the **user login script** if it exists in the user's home (login) directory:

If the Shell is ...	The User Login Script is ...
Korn (/bin/ksh)	.profile
Bourne (/bin/sh)	.profile
Restricted (/bin/rsh,/bin/krsh)	.profile
C (/bin/csh)	.login

Every user will have the appropriate local login script in his/her home directory already. It was copied there by SAM when you added the user. Users can customize their environments by modifying these files to suit their needs.

Typical changes to these files include:

- Changing PATH so the shell searches different directories for commands.
- Setting a specific value for TERM to match the user's terminal type.
- Setting other user specific environment variables.

Note

If all users on your system use the same applications then it is easiest to set up any necessary environment variables in the system login scripts.

If a few users use different applications or require different environment settings these are most easily handled in the local login script.

5-9. SLIDE: Customizing the User Environment Instructor Notes

Purpose

To show the two different ways an administrator can set up the correct user environment.

Be sure that students understand the differences between the "system login script" and the "user login script".

Teaching Tips

- This is important because almost every application requires changes to the user environment.
- This can also be very difficult to understand at first. Try to keep it as simple as possible. We all know there are many things about the user environment and its setup that are not given here. Just remember your audience and give them only what they need.

Default Local Login Files

The following files are copied to the new user's home directory by SAM when the user is added to the system:

- `/etc/d.profile` - "ksh" and "sh" local login script
- `/etc/d.login` - "csh" local login script
- `/etc/d.cshrc` - "csh" startup script (run every time a csh is started).
- `/etc/d.exrc` - vi configuration file

Transition

Let's look at an example of customizing a user environment.

5-10. SLIDE: A Sample Environment Setup

A Sample Environment Setup

All users need `/usr/widget/bin` in their PATH and `vi` for their EDITOR so change `/etc/profile`:

```
PATH=/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:/usr/widget/bin
EDITOR=vi
export EDITOR
```

Norton prefers `emacs` for EDITOR so he can change his `.profile`:

```
EDITOR=emacs
```

50722 5-10 56

Student Notes

In this example, let's assume there are two users accounts on the system. Their user names are *ralph* and *norton*. Both users have the Korn shell (`/bin/ksh`) as their startup program.

ralph and norton use a fictitious application called "Widget". Widget requires some changes to the user environment:

- The directory `/usr/widget/bin` must be part of the PATH.
- The EDITOR variable must be set to `vi` or `emacs` depending on the user's preference for a text editor.

System Login Script Changes

Since both ralph and norton use Widget they require the same basic environment configuration. You can set up this default environment by editing the appropriate lines of the `/etc/profile` file.

The PATH variable setting looks something like this by default:

```
PATH=/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin  # default path.
```

The new PATH setting for the Widget application should look something like this:

```
PATH=/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:/usr/widget/bin
```

You could also set up the default editor of choice to be `vi` in `/etc/profile` by adding these lines:

```
EDITOR=vi
export EDITOR
```

Now when ralph and norton next login they will both be able to use the Widget application correctly, and if they invoke the editor in the application, they will get `vi`.

User Login Script Changes

If norton prefers to use the `emacs` editor instead of the default, he (or you) could modify the `.profile` file in norton's home directory. The lines to set the EDITOR variable to emacs might look like this:

```
EDITOR=emacs
```

Note that EDITOR does not need to be exported because it was already exported in `/etc/profile`.

The next time norton logs in his EDITOR variable will be set to "emacs". When norton invokes the text editor from the Widget application he will execute the `emacs` editor instead of `vi`. Note that ralph is unaffected and will still use the `vi` editor.

5-10. SLIDE: A Sample Environment Setup — Instructor Notes

Purpose

Customizing the user environment can be complicated. This is a simple example to show where and how changes might be made for a specific application.

Key Points

- PATH is one of the most commonly changed environment variables.
- TERM is also a very import variable and will be discussed when we describe terminals.
- vi and emacs are by far the most commonly used text editors on a UNIX system. That's why they are used here as examples.
- Remind the students that the login scripts are only run at login. If a change is made while a user is logged in it will not take effect until the next time they login.

Teaching Tips

- If you know anything about the student's application and the environment setup, please relate that to the class. Everyone can benefit from real-life examples.
- Depending on the level of knowledge your students have you may want to show how environment variables can be set and changed interactively from the shell command line.
- Another common example of an environment variable that may be used is the K-shell's `ENV` variable. `ENV` contains the name of a file to be executed every time a K-shell is started (when you start a shell from within SAM, for example). This file usually contains things like alias commands and prompt changes. The value of ENV is typically set in the user's **`.profile`**:

```
ENV=$HOME/.kshrc
export ENV
```

- The K-shell has an alternate way to set and export variables simultaneously:

```
export EDITOR=emacs
```

Transition

Everything we have discussed concerning the user environment has assumed that you are logging in from a alphanumeric terminal or a generic X windows display. If you are using HP-VUE the login files are sightly different.

5-11. SLIDE: The HP-VUE User Environment

The HP-VUE User Environment

Login files are **NOT** the same in HP-VUE

- Many environment variables are built in:
  DISPLAY
  USER
  HOME
  PATH
  SHELL
  TZ
- System login file (used instead of `/etc/profile`):

  `/usr/lib/X11/vue/Vuelogin/Xsession`

- User's local login file (used instead of `.profile`):

  `.vueprofile`

50722 5-11 57

Student Notes

Logging in on a bit-mapped display with the HP Visual User Environment (HP-VUE or VUE) user interface is different than login on an alphanumeric terminal. The software that controls your login session is more complex than the simple shell that normally controls a terminal session.

Because of these differences, the setup of the user environment is also somewhat different. The normal shell startup files (like `/etc/profile` and `.profile`) are no longer used. Instead, certain environment variables are built into VUE and others can be set or modified with two files specific to VUE.

The environment variables that are built into VUE are:

Variable	**Value**
DISPLAY	Set to the name of the X windows display. The default is `system_hostname:0`.
USER	User name.
HOME	Home directory.
PATH	Defaults to `/usr/bin/X11:/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin` .
SHELL	The user's startup shell
TZ	The system time zone specifier

System Login File

The VUE environment uses `/usr/lib/X11/vue/Vuelogin/Xsession` to set up the default user environment similar to how you can set up the default environment in `/etc/profile`. The variables set in `Xsession` and their defaults are:

```
EDITOR=/usr/bin/vi
LOGNAME=$USER
TERM=hpterm
MAIL=/usr/mail/$USER
```

Other system-wide environment variables needed for your applications can also be placed in the `Xsession` file.

User Login File

User's can control their specific environment by editing `.vueprofile` which must reside in their home directory. This file is used very much like the user's `.profile`. A default copy may be found in `/usr/lib/X11/vue/sys.vueprofile`.

The first line of `.vueprofile` must name the shell that is specified with the user account. For example,

If the user's shell is . . .	**The first line should be . . .**
Korn (`/bin/ksh`)	`#!/bin/ksh`
Bourne (`/bin/sh`)	`#!/bin/sh`
C (`/bin/csh`)	`#!/bin/csh`

It is important that the `#!` be the very first characters in the file. The `.vueprofile` file should contain only lines that set environment variables. It must **NOT** contain commands that do terminal control (like `tset` and `stty` which are used in `.profile`).

The Widget Example in VUE

If we wanted to extend our example of ralph and norton to the VUE environment we would have to edit `/usr/lib/X11/vue/Vuelogin/Xsession` and `/users/norton/.vueprofile`.

The `Xsession` file would contain the default PATH setup for both ralph and norton. The line would look something like this:

```
PATH=/usr/bin/X11:/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:/usr/widget/bin
```

To allow norton to use the `emacs` editor as we did before, norton's `.vueprofile` should look like this (recall that norton uses the Korn shell when he logs in):

```
#!/bin/ksh
EDITOR=emacs
export EDITOR
```

Reference

The *HP-VUE System Administration Manual* has more details on using and configuring the VUE environment.

The *HP-VUE User's Guide* is also an excellent introduction to using and customizing HP-VUE.

5-11. SLIDE: The HP-VUE User Environment — Instructor Notes

Purpose

The VUE environment is set up differently than the standard shell (terminal) login environment. This is a very brief introduction to the very complex task of administering a custom VUE environment.

Key Points

- `/etc/profile` and `.profile` are not used in VUE except if you start a login shell window (`hpterm -ls` will do this).
- Environment settings in `Xconfig`, `Xsession`, and `.vueprofile` will apply to all shell windows and applications that run in the VUE environment.
- The K-shell ENV variable can be set the same way in `.vueprofile` as it is in `.profile`:

  ```
  ENV=$HOME/.kshrc
  export ENV
  ```

  Now every window started in the VUE session that runs a K-shell will run the user's .kshrc file. This is very important to VUE users that are used to having a .kshrc file.

Teaching Tips

- Most of these users will not want to know the details of VUE and the many resources you can set. They just want to get their applications working in their environment.
- If they (or you) want more information on configuration of the VUE environment see the *HP-VUE System Administration Manual.*

5-12. WORKSESSION: Review Questions

1. Describe the user account information needed to build a user account on the system.

2. What is the difference between removing a user and deactivating a user on the system?

3. List three ways a user's password may be changed.

4. Why is it important to protect directories (especially user home directories) with 755 permissions?

5. What is the difference between `/etc/profile` and `.profile`?

6. If you were planning to shutdown the system at 8:00pm, how might you communicate that to all of your users?

7. When 8:00pm arrives, how can you immediately warn anyone still logged in that the system will be shutdown now?

5-12. WORKSESSION: Review Questions

Instructor Notes

These review questions are very quick questions that make the students revisit the material and reinforce the important topics.

The topic references are provided when possible to more closely tie into the material.

1. Describe the user account information needed to build a user account on the system.

Answer:

See topic "Adding a New User with SAM" for details.

This information is stored in `/etc/passwd` in the following format:
`user_name:password:user_id:group_id:comments:home_dir:startup_program`

2. What is the difference between removing a user and deactivating a user on the system?

Answer:

Reference: "Removing User Accounts" and "Deactivating and Reactivating a User"

Removing a user actually takes the account off the system permanently by removing the user's entry from the `/etc/passwd` file. The only way to get the user back on the system is to add the account again.

Deactivating a user simply changes the user's password to *. The user's account still exists but s/he can not log in.

Both removing and deactivating a user have several actions that may be taken regarding the user's files.

3. List three ways a user's password may be changed.

Answer:

Reference: "Setting a User Password" under "Changing the User's Password"

1. Using SAM's "Change a User Password" option.
2. By the superuser using the `passwd user_name` command.
3. By the user using the `passwd` command.

4. Why is it important to protect directories (especially user home directories) with 755 permissions?

Answer:

Reference: "User Security Considerations"

755 permissions look like this on a directory:

```
drwxr-xr-x  2 ralph    users       1024 Jun 17 15:00 /users/ralph
```

Notice that anyone can list the directory contents (the r permissions) and anyone can cd to, or otherwise use, the directory (the x permissions) but only ralph (the owner) can change its contents (the w permission).

This prevents everyone else from:

removing files from the directory (`rm` looks at directory permissions)
creating files in the directory (like network security files)
renaming files in the directory

5. What is the difference between `/etc/profile` and `.profile`?

Answer:

Reference: "Customizing the User Environment"

`/etc/profile` is the system login file used for all users. Anything that the system administrator puts in this file will have an effect on all users (running `ksh`, `sh` or a restricted shell) when they log in.

The `.profile` file is the user's own file. Any command that s/he adds to this file will be executed when s/he logs in and only when this user logs in, no one else. This is how a user can make his/her environment different from the defaults set in `/etc/profile`.

6. If you were planning to shutdown the system at 8:00pm, how might you communicate that to all of your users?

Answer:

Reference: "Communicating with System Users"

Using the `/etc/motd` file. Anything that you put into the `/etc/motd` file will be displayed whenever a user logs into the system.

7. When 8:00pm arrives, how can you immediately warn anyone still logged in that the system will be shutdown now?

Answer:

Reference: "Communicating with System Users"

Using the `wall` command. This will send a message to everyone that is logged into the system.

5-13. LAB: Exercises

Directions

Only one person at a time (per system) should use SAM. If more than one person or lab group is logged into the same machine you must coordinate your activities somehow.

For series 800 students: if the pseudo superuser environment is available, each pseudo superuser group can do the labs simultaneously. Ask your instructor if and how the pseudo superuser accounts are set up on the Series 800 used for class.

1. When it is your turn, invoke `sam` to add a new user account to your system. Remember that you must be logged in as `root` (the superuser) to invoke SAM.

Give the user the following attributes:

- Use your (or your partner's) name as a user name.
- Use the default group and home directory.
- Use the default (Bourne) shell as the startup program (`/bin/sh`).
- Use any comments you would like.
- Don't modify the user defaults.
- Make the user assign their own password when they first log in.

2. Now, exit SAM and look at the `/etc/passwd` and `/etc/group` files. Do you see the user you added?

3. Log out as `root` and log back in using your new account. Try changing the password on your account to your first name. Does this work ? Why or why not? Set a valid password on the account so you can complete the login.

4. Once logged in, execute `ll -a` to view the files that were created for the user.

Notice the permissions on the files and directories.

5. While logged in as your user, execute the `env` command and note some of the important things in the environment line `PATH`, `TERM`, `LOGNAME`, etc.

6. This is an exercise to demonstrate the importance of having the user environment set correctly.

Execute `sh` to create a new shell we can experiment with. Now execute `unset TERM` this removes that value that describes your terminal type. Commands like `vi` will not work correctly without a valid value in `TERM`. Execute `vi afile`. What happens? Why?

In `vi` type `:q!` to exit back to the shell then type Ctrl d to exit the extra shell we started.

7. Log out as your user and log back in as `root`. Using SAM, deactivate your user.

What happens if you try to log in as the user now?

Log back in as `root` and reactivate the user with SAM.

8. Create a second user account using whatever user information you want.

When the user has been added use SAM's `Shell` softkey to start a shell. Now look in `/etc/passwd` and the user's home directory to convince yourself that the user really exists. Exit the shell by typing `exit`.

Now remove the user and all of his/her files from the system. Once the task completes read the message displayed by SAM.

You can later check `/etc/passwd` and the user's home directory and notice that the account is gone.

9. Advanced:

Use SAM to change the shell your user use at login to the Korn shell (`/bin/ksh`).

5-13. LAB: Exercises

Instructor Notes

1. When it is your turn, invoke sam to add a new user account to your system. Remember that you must be logged in as root (the superuser) to invoke SAM.

Give the user the following attributes:

- Use your (or your partner's) name as a user name.
- Use the default group and home directory.
- Use the default (Bourne) shell as the startup program (/bin/sh).
- Use any comments you would like.
- Don't modify the user defaults.
- Make the user assign their own password when they first log in.

Answer:

```
# sam (Return)

    Users ->

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

Fill in the information a requested on the data entry screen and press the **Perform Task** softkey.

When SAM prompts you for a password, enter ,.. to force the user to set a password the first time they log in.

2. Now, exit SAM and look at the /etc/passwd and /etc/group files. Do you see the user you added?

Answer:

To exit SAM use the **Main Menu** softkey then **Exit**.

To see the contents of the files use the more command:

```
# more /etc/passwd (Return)
```
Contents of /etc/passwd

```
# more /etc/group (Return)
```
Contents of /etc/group

In /etc/passwd you should see a line that starts with the new user's name and contains all of the information you used on the data entry screen. Notice the ",.." that you entered for a password in the second field. If you had assigned a real password to this user it would show up here in an encrypted form.

In /etc/group you will find a line that starts with the group name users and your new user will be listed as a group member.

3. Log out as root and log back in using your new account. Try changing the password on your account to your first name. Does this work ? Why or why not? Set a valid password on the account so you can complete the login.

Answer:

Since we set up the user to set a password at first login, you will enter the new user's name at the login: prompt and immediately be prompted for a new password.

You will not be able to change it to your name because of the password naming restrictions insist on at least one non-alphabetic character.

4. Once logged in, execute ll -a to view the files that were created for the user.

Notice the permissions on the files and directories.

Answer:

The files and permissions should look something like this:

```
$ ll -a
total 12
drwxr-xr-x   2 ralph    users       1024 Jun 17 10:33 .
drwxr-xr-x  55 root     sys         1024 Jun 17 10:33 ..
-rw-r--r--   1 ralph    users        818 Jun 17 10:33 .cshrc
-rw-r--r--   1 ralph    users        347 Jun 17 10:33 .exrc
-rw-r--r--   1 ralph    users        377 Jun 17 10:33 .login
-rw-r--r--   1 ralph    users        382 Jun 17 10:33 .profile
$
```

Note that only the user is allowed to modify the home directory (.) and the login files.

5. While logged in as your user, execute the env command and note some of the important things in the environment line PATH, TERM, LOGNAME, etc.

Answer:

The env command displays your entire working environment. A default user environment will be similar to this:

```
$ env
EDITOR=vi
HOME=/users/ralph
LOGNAME=ralph
MANPATH=/usr/man:/usr/contrib/man:/usr/local/man
PATH=/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:.
```

```
SHELL=/bin/sh
TERM=hp
TZ=EST5EDT
$
```

6. This is an exercise to demonstrate the importance of having the user environment set correctly.

Execute `sh` to create a new shell we can experiment with. Now execute `unset TERM` this removes that value that describes your terminal type. Commands like `vi` will not work correctly without a valid value in `TERM`. Execute `vi afile`. What happens? Why?

In `vi` type `:q!` to exit back to the shell then type (Ctrl)(d) to exit the extra shell we started.

Answer:

```
# sh (Return)
#
```
This is the new shell prompt
```
# vi afile (Return)
I don't know what kind of terminal you are on - all I have is 'unknown'.
[Using open mode]
"afile" [New file]
:q!
#
# exit (Return)
#
```
Now we are back to the original shell

As `vi` doesn't know anything about the terminal capabilities (like number of lines or characters per line) without the TERM environment variable, it assumes you are using a line oriented teletype. Though the appearance in `vi`'s open mode is unfamiliar, we still have the complete `vi` command set at our disposal. The best way out is :q! and a checking and re-setting of the TERM variable.

7. Log out as your user and log back in as `root`. Using SAM, deactivate your user.

What happens if you try to log in as the user now?

Log back in as `root` and reactivate the user with SAM.

Answer:

```
# sam (Return)

    Users ->

        Deactivate a User Account ...
```

When you try to log in as the deactivated user, you will always get the `login incorrect` message. This is because the password in `/etc/passwd` has been changed to `*`. This is a special password that will never match anything the user types.

To reactivate the user select `Reactivate a User Account ...` .

8. Create a second user account using whatever user information you want.

When the user has been added use SAM's `Shell` softkey to start a shell. Now look in `/etc/passwd` and the user's home directory to convince yourself that the user really exists. Exit the shell by typing `exit`.

Now remove the user and all of his/her files from the system. Once the task completes read the message displayed by SAM.

You can later check `/etc/passwd` and the user's home directory and notice that the account is gone.

Answer:

Adding a user is identical to the earlier exercise.

To remove the user:

```
# sam (Return)

    Users ->

        Remove a User Account from the System ...
```

Fill in the user name and specify that you want the user's files removed also.

SAM will complete the task and notify you that the user account is gone but the disk search to remove all of the user's files may take some time. This means the entry from `/etc/passwd` is taken out but it may be a few minutes before all of the user's files are removed.

9. Advanced:

Use SAM to change the shell your user use at login to the Korn shell (`/bin/ksh`).

Answer:

```
# sam (Return)

    Users ->

  View/Modify a User's Account Information ...
```

On the data entry screen, move down to the "Start-up Program" field and type /bin/ksh or use the Help softkey to list the shells so you can select one. When you press Perform Task the shell will be changed.

To test this you can log in as the user and execute env. The SHELL variable will be set to /bin/ksh. You could also execute the ps (process status) command and you would see that you are running ksh.

Module 6 — Printers and Plotters

Objectives

Upon completion of this module, you will be able to do the following:

- Add a printer to your system.
- Start and stop the printer (LP) spooler.
- Remove a printer from your system.
- Check the spooler status.
- Enable and disable a printer.
- Set the default printer.
- Move print requests to other destinations.
- Cancel print requests.

About this Module

This module will discuss the LP (line printer) spooler system that is available with HP-UX. In this module we will look at the terminology of the spooling system, commands that users can invoke to interact with the system, and administrative commands and tools to monitor the status of the spooler. In particular, we'll look at how printers are configured for the system and how a printer can be designated as the system default printer.

We will use SAM for most of the tasks involved with configuring the spooling system.

Any references to printers in this module also apply to plotters that are controlled through the LP Spooling system.

Overview of Module 6

Printer administration is a task that almost everyone in the class will undertake at some time.

We will be using SAM for almost everything in this module. A few commands like `lpstat` are shown as convenient ways of doing tasks without SAM. Commands like `cancel` and `lpmove` are covered because they cannot be done through SAM.

Network printer setup is covered in the instructor notes but not in the student notes.

6-1. SLIDE: Spooling System Plumbing Diagram

Spooling System Plumbing Diagram

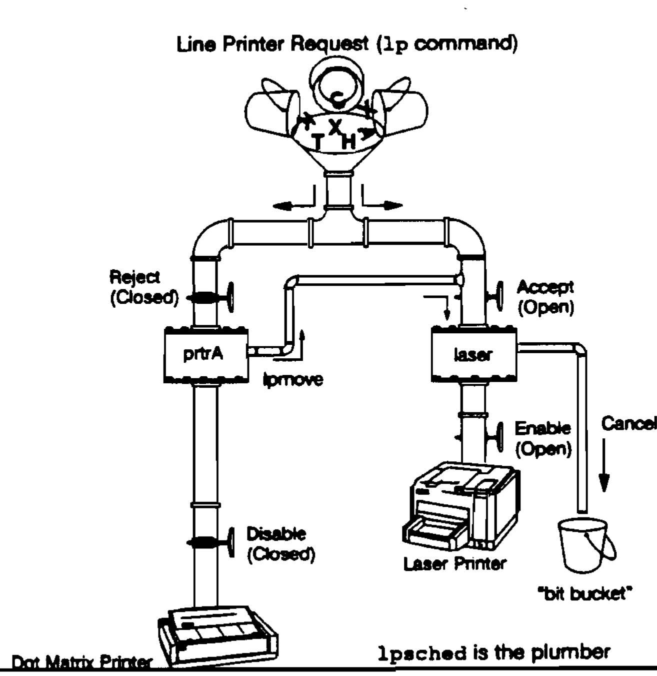

50722 6-1 58 ©1992 Hewlett-Packard Co.

Student Notes

On a multi-user system, access to the printers requires careful management and control. Since users could possibly send print requests to the same printer at the same time, there needs to be a way to:

- Make sure each request is printed separately.
- Determine which request will be printed first.

Definition

The **LP spooler** is a collection of utilities and commands that controls the print requests of users. The LP spooler stores print requests in a spool directory until a printer is available. When a printer is available, print requests are processed, one at a time, in the background.

The spooling system must be configured and customized by the administrator. For example, the administrator will add printers to the spooling system, then choose where requests are spooled by default, move requests between printers, and decide which printers will accept requests.

After the LP spooler system is installed, any user can submit a job to be printed, obtain the status of all or any printer, cancel any print job, or declare printers to be in and out of service.

To submit a job to the spooling system any user will use the `lp` command. The basic form of the `lp` command looks like this:

`lp [`*options*`]` *filename* `[`*filename* . . .`]`

- *filename* — is the pathname of the file or files to be printed.
- *options* — is the list of possible modifiers to the basic print request. Some of the most common options include:

-d*printer_name*	Spools the request to *printer_name* instead of the default printer. *printer_name* is the name you gave the printer when you added it to the system using SAM.
-n*number*	Make *number* copies of the requested file(s).
-t*title*	Print *title* in large letters on the header page of the printout.
-o*printer_specific_option*	This modifies the request according to the *printer_specific_option*. The names of these options can be found in the top of the file `/usr/spool/lp/interface/`*printer_name*. Some common options include:

-onb	Do not print a banner page.
-oraw	The request is graphics not text.
-ol*lines*	Set page length to *lines* lines.
-ohalf	Print two pages on one side of a sheet (LaserJets only).

For example, the command line:

```
# lp -onb -n2 -dlaser filex filey
```

will print 2 copies each of "`filex`" and "`filey`" to the printer called "laser" without printing a header page.

Reference

The *System Administration Tasks* manual has a good chapter on the LP spooler system called "Managing Printer Output".

The details of the `lp` command can be found in section 1 of the *HP-UX Reference Manual*.

6-1. SLIDE: Spooling System Plumbing Diagram Instructor Notes

Purpose

This is a broad overview of the concepts of the LP spooler and some of the administration tasks involved. It is also to give a quick review of `lp` since you can't very well test a spooling system without it.

Key Points

- The thought of any user having the ability to cancel any print job strikes most administrators as dangerous. Assure them that if someone cancels a job other than their own, the user who's job was canceled is notified by mail that their job was canceled and who did it!
- If an administrator doesn't want users to be able to disable printers, simply change the permissions on `/usr/bin/disable` so non-root users can not execute it.

OpenSpool/UX

OpenSpool/UX is an HP product to increase administrators' productivity in managing their spooling environment, whether it be a standalone or a networked environment. This new tool is offered on *all* HP 9000 systems. OpenSpool/UX was developed with multivendor cooperation in mind. Users have their choice of how to interface with OpenSpool/UX: command-based, screen interface or Motif-windows.

Teaching Tips

- Discuss the flow of a print request on the slide. This plumbing diagram is referred to several times throughout the module.
- Point out that any user can cancel requests, disable and enable printers.
- Point out that accept, reject and lpmove are administrator only commands.
- Some analytical types will wonder why we don't show lines to move requests between printers in both directions (why aren't there two lpmove pipes?). The answer is simple: There is not enough room on the slide to show all possible combinations of functions!
- You may want to mention that there is a "plumber" process that controls all of this called `lpsched`.

Preparation Notes

The *System Administration Tasks* manual has a good chapter on the LP spooler system called "Managing Printer Output".

Transition

Before we can do anything else we need to add a printer to our system.

6-2. SLIDE: Managing Printers with SAM

Managing Printers with SAM

```
SAM                          Printers and Plotters

    The choices below pertain to plotters as well as printers. Highlight an item and
    then press "Return" or "Select Item".

          Printers:
                  Add a Local Printer ...
                  Add a Remote Printer ...
                  Enable a Printer ...
                  Disable a Printer ...
                  Remove a Printer ...

          LP Spooler Administration:
                  Set the System Default Printer ...
                  Shut Down the LP Spooler
                  View Printer Status Information
                  View Print Requests

Help | Main Menu | Shell | Select Item | hpterm | | | | Previous Menu
```

50722 6-2 59

Student Notes

SAM can be to do most of the administration of the the LP spooler system. The tasks that SAM can perform are shown on the slide.

```
# sam (Return)

    Peripheral Devices ->

  Printers and Plotters ->
```

There are some commands that are not available in SAM. For example, SAM cannot:

- submit a request
- move a request between printers
- cancel a request
- reject requests without disabling the printer

6-2. SLIDE: Managing Printers with SAM

Instructor Notes

Purpose

To point out what SAM can and cannot do when administering the LP spooler.

Key Points

- There are two ways of performing the spooling system administration, SAM or a set of individual commands. SAM is easier to use, the individual commands more flexible. For this audience we will stick mostly to what SAM can do.

Teaching Tips

- Only cover the slide in general terms. The details will come throughout the rest of the chapter.
- Point out the things SAM cannot do. For these tasks we will need to use the commands.

Transition

6-3. SLIDE: Adding a Printer with SAM

Adding a Printer with SAM

```
SAM                         Add a Local Printer

    Fill in or modify the desired fields and then press "Perform Task".

    Printer name . . . . . . . . . . . . . . . . .  laser
    Printer model/interface  . . . . . . . . . . .  laserjet
    Printer device file name . . . . . . . . . . .  /dev/lp_laser
    Printer priority (0 - lowest, 7 - highest)  . . 0
    Make this the system default printer? (y or n)  y
    Printer class . . . . . . . . . . . . . . . . .  ____________ (optional)
    Printer connected to a terminal running TSM? (y or n)  n  (optional)

 Help | Main Menu | Shell | Perform Task | hpterm |  |  |  | Exit Task
```

50722 6-3 60

Student Notes

To add a printer with SAM, choose the following and you will get the screen pictured on the slide:

```
# sam (Return)

     Peripheral Devices ->

         Printers and Plotters ->

             Add a Local Printer ...
```

Note

SERIES 800 NOTE: On the Series 800 you must also select which type interface the printer being added will use (serial, HP-IB, Parallel or connected to a terminal running TSM). Once you have provided the printer type, SAM will display all of the available interfaces of that type on your system. Simply select the correct interface (and the port number if it is a serial device) then SAM will display the screen pictured on the slide.

Now you must fill in the information regarding your new printer. The screens may vary slightly in format on your system but all of the vital information is the same:

- "Printer Name" — This is the name that will be recognized by the LP Spooler system. This is the name you will use to submit requests to the printer and configure any aspect of this printer's queue.
- "Printer Model/interface" — This is the type of printer you are adding. This name actually specifies the program that will do all of the printer specific setup according to the options you use with the `lp` command. If you are not sure what model to use, press the `Help` softkey. SAM will give you a list of the supported printers.
- "Printer device file name" — This is the device file that references the printer. If the device file does not exist, SAM will create it for you. On a Series 300/400/700 this name is `/dev/lp_`*Printer_Name* by default. On a Series 800 the default name depends on the printer's interface type. A serial printer will have a device file name of `/dev/tty`*LpP*. Other types of printers will be called `/dev/lp`*L*. (Where *L* is the logical unit number of the printer or MUX and *P* is the port number of the MUX.)
- "Printer Priority" — This is the default priority of print requests coming into this printer's queue. A user can submit requests with higher or lower priorities than this number using the `lp -p`*priority* command. Requests with the highest priority in the queue will print first. If all requests have the same priority they print in a first-in-first-out (FIFO) order.
- "Make this the default system printer" — If you answer y this will become the printer that all requests go to unless they were submitted using the `lp -d`*printer_name* command.
- "Printer Class" — You can add a group of like printers to a class. Any request submitted for the class will go to the first available printer in the class. This works well only if you have several of the same type of printers in the same general area. You can not control which printer in the class will get your request.
- "Printer connected to a terminal running TSM" — TSM is the Terminal Session Manager. It allows a user to run multiple sessions on one alphanumeric terminal display. If you connect a printer to the serial "printer" port on the back of a terminal running TSM, SAM can configure the spooler to print to that printer without disrupting the user that is logged in there. This question does not appear on the Series 800 screen because you would select this option from the previous menu.

Press the `Perform Task` softkey and SAM start to add the printer to the spooling system.

Note

SERIES 300/400/700 NOTE: SAM will now prompt for device file information if it does note exist. You will have to first specify the type of printer (serial, parallel, HP-IB, etc.) and on the Series 300/400 the select code of that interface. You may also have to specify the port number or bus address (depending on the interface type). If you have any questions about the interface types or other numbers, simply press `Help` and SAM will list your choices.

Once the task has completed, SAM will ask you if you want to test the device. If you answer y, a small test file will be printed.

Example

On the slide we are adding a LaserJet printer to the spooling system with the following characteristics:

- The printer's name will be `laser`.
- It will use the model/interface program for LaserJet printers called `laserjet`.
- Its device file name will be the default `/dev/lp_laser` which SAM will create for us if needed.
- Its request priority will be 0.
- It will be the system default printer.
- It will not be part of a printer class.
- It will not be connected to a terminal running TSM.

Reference

The *HP-UX System Administration Tasks* manual has a few more details on what SAM is prompting for.

You can find more information about printer priorities and using them for setting up batch style printing by looking up the `lpfence` command in the *HP-UX Reference Manual*.

The *HP-UX Reference Manual* also has more information about TSM in section 1.

6-3. SLIDE: Adding a Printer with SAM — Instructor Notes

Purpose

To explain how to add a printer to the LP Spooler system using SAM. This is all that really needs to be done to make a printer fully functional in SAM.

Key Points

- SAM is much easier for adding a new printer to the spooling system than using the corresponding `lpadmin` command.
- SAM will create the appropriate device file for you if it does not exist.
- Printer classes don't make sense if the printers are scattered throughout the facility because you will not know where to find your output. Few people use classes today.

Teaching Tips

- Go over the example on the slide: On the slide we are adding an HP LaserJet III printer (with a parallel interface) to the spooling system. We named the printer "laser". For the model/interface we used `Help` and selected "laserjet". We can accept the default device file name. We want the printer to be our default printer. We do not want it to be part of a class so the field is left blank. We cannot connect it to a terminal running TSM because we want it connected using the parallel interface (a terminal's printer port is always serial).
- You may want to point out the details of the popup windows that prompt for device file information. Note how they are different for 3/400s vs. 700s vs. 800s. Note also that 3/4/700s ask for device file information after you press **`Perform Task`** whereas the 800s use menus to ask for this information first.

A Note on TSM

TSM is the Terminal Session Manager (it comes with the core software set but loading it is optional). This package allows a user of an alphanumeric terminal to run multiple sessions from the same terminal simultaneously. You can switch between the sessions using function keys, menus or via a command line interface. One other TSM session is dedicated to the printer port on the back of HP terminals. If you connect a printer to this port and have SAM set this printer up as a spooled printer, you can use the printer without affecting the other sessions that are running.

Transition

Now that we have a printer configured we will want to see the status of the spooling system and what requests are queued.

Handout: Adding a Remote Printer

We are not going to cover adding a remote printer as part of the class but you may choose to do this as an added topic if the network is set up in the classroom. If you do not have a printer for every system, setting up network printer access will allow all of the students to do most of the exercises to the remote printer.

Adding A Remote Printer with SAM

A remote printer is a printer connected to another system on the network where you would like to send print requests. The system physically connected to and controlling the printer is called the **printer server**. Any systems that want to print to this printer from across the network are called **printer clients**. The printer server will be configured into the client's spooling system as another printer queue and is controlled just like another local printer.

Adding a remote printer using SAM is very similar to adding a local printer. Before you can configure the client systems using SAM, however, you must set up the server.

To configure the *printer server* you must:

1. Add the printer to be shared as a local printer as normal.
2. Edit `/etc/inetd.conf` to start the printer server daemon `rlpdaemon`. Find the line that looks like this:

```
#printer     stream tcp nowait root /usr/lib/rlpdaemon  rlpdaemon -i
```

   and remove the leading `#`.

3. Reconfigure the `inetd` networking daemon so it will automatically start the `rlpdaemon` when needed:

```
# inetd -c
```

4. Make sure all clients are listed in the server's `/etc/hosts` file. This should have already been done as part of the network setup.

To configure the *printer clients* you must:

1. Setup/verify network communications between the client and server systems.
2. Use SAM to `Add a remote printer ....`
   - Give the printer a local name so the local users can access it. This does not have to be the same name the printer server gave to the printer but it typically is the same to avoid confusion.
   - Specify the printer server's network hostname.
   - Specify the name of the printer on the server.
   - Use the default values for the "Remote cancel model" and "Remote status model" fields.
   - Specify if this will be the client's default printer.
   - Fill in optional fields if needed.

Once the client and server systems are configured the users on the client system can access the remote printer on the server as if it were directly connected to the client machine.

6-4. SLIDE: Checking Printer Status

Checking Printer Status

- View Printer Status Information

  What is the default printer name?

  Is the spooling system active?

  Which printers are enabled, disabled or printing?

- View Print Requests

  What requests are pending in all printer queues?

50722 6-4 61

Student Notes

The slide describes exactly what SAM can do to allow you to see the printer and spooling system status.

There is one command outside of SAM that is extremely useful and you will use it often. It is the `lpstat` command.

`lpstat -t` allows you to check the status of all parts of the spooling system. It prints the status of the scheduler and each printer configured in the spooler. For each printer, it tells you if that printer is accepting or rejecting requests, if it is enabled or disabled, the default priority and the request-id of the requests that are queued for printing.

```
# lpstat -t
scheduler is running
system default destination: laser
device for laser: /dev/lp_laser
device for dustylp: /dev/tty0p2
laser accepting requests since Jun 14 15:06
dustylp accepting requests since Apr  2 13:06
printer dustylp is idle.  enabled since Apr  2 14:32
     fence priority : 0
printer laser now printing laser-1140.  enabled since Jun 17 11:39
     fence priority : 0
laser-1140          ralph          priority 0  Jun 24 09:24 on laser
     /etc/group                2    copies   214 bytes
```

Reference

There are many more options to `lpstat`. The *HP-UX Reference Manual* lists them all.

6-4. SLIDE: Checking Printer Status

Instructor Notes

Purpose

We need a way to see what our printers are doing if we are to control them with SAM or the commands on the next slide. `lpstat -t` gives you all of the information necessary to intelligently manage the printers.

Key Points

- If you are not currently in SAM, it is easiest to find spooler status information using the `lpstat` command.

Teaching Tips

- You may want to mention the `-p` and `-o` options of `lpstat`. If students have many printers, the `lpstat -t` output can become too long to read easily. These options help reduce the size of the output so you can see just a few printers.

  The `-p` option lists only the status of the specified printers and the `-o` option lists the requests queued for the specified printers.

Transition

Now that we have at least one printer we should look at some of the things to do to manage

6-5. SLIDE: Controlling the Spooling System

Controlling the Spooling System

- `Enable a Printer ...`   `Disable a Printer ...`

  Allow/disallow a printer to accept requests and print them.

- `Start Up the LP Spooler`   `Shut Down the LP Spooler`

  Allow/disallow *all* printer requests to print.

- `cancel`

  An HP-UX command that cancels print requests.

- `lpmove`

  An HP-UX command that moves print requests between printer queues.

50722 6-5 62

Student Notes

As we saw in the "Plumbing Diagram" of the spooling system, there are valves above and below the printer queue (tank). If you select `Disable a Printer ...`, both of these valves are closed for the specified printer. This means that requests already in the queue will stay in the queue and will not print. If anyone attempts to submit a new request with `lp` they will get an error indicating that the printer is not accepting any more requests. Selecting `Enable a Printer ...` opens both of these valves.

To stop all printers from printing their queued requests you can `Shut Down the LP Spooler ...`. When the spooler is shut down all printers stop printing. Any requests that were partially printed will start again from the beginning when the spooler is restarted. Note that while the spooling system is shutdown, jobs can still be submitted to printers using the `lp` command. Of course, `Start Up the LP Spooler ...` restarts all printing.

There are also ways to accomplish these tasks from the command line:

- `reject [-r` *reason*`]` *printer_name* — This will keep *printer_name* from accepting any further requests. Users that try to submit requests using `lp` will get an error. If you specified the `-r` *reason* option, the *reason* will be given as part of the error.
- `accept` *printer_name* — Allows *printer_name* to accept new requests. `accept` and `reject` control the upper valve of the spooling system.
- `disable [-r` *reason*`]` *printer_name* — This will keep *printer_name* from printing any requests in its queue. If you specify the `-r` *reason* option, the *reason* for being disabled will be displayed with the printer status.
- `enable` *printer_name* — Allows *printer_name* to print queued requests. `enable` and `disable` control the lower valve of the spooling system.
- `lpsched` — This is the "scheduler" daemon that controls all printing. No printing will occur unless `lpsched` is running. Shutting down the spooler kills this process.
- `lpshut` — This is the program that kills `lpsched`.

The other command on the slide is `cancel`. This is the only way to cancel requests that are queued for a printer. `cancel` can cancel requests in one of three ways:

- `cancel` *request-id* — Cancel just the specified request. The *request-id* is returned by `lp` and you can see it using `lpstat`.
- `cancel` *printer_name* — Cancel the request that is currently printing on *printer_name*.
- `cancel` *printer_name* -e'' — Cancel all requests queued for *printer_name*.

Note

Anyone can `cancel` a print request even if it is not their own. The user will be notified, however, by an electronic mail message that describes the request that was canceled and who canceled it.

Only the superuser can use the `-e` option.

The `lpmove` command can only be executed by the superuser. It is used to move one or all requests from one print queue to another.

`lpmove` *request-id printername*

— Move the request called *request-id* to *printernames* queue.

`lpmove` *printer1 printer2*

— Move *all* requests from *printer1*'s queue to *printer2*.

6-5. SLIDE: Controlling the Spooling System

Instructor Notes

Purpose

To describe the ways to control the queue with SAM and to briefly mention the manual commands.

Key Points

- The startup and shutdown options do not actually control the enable/disable valves on the print queue. Instead they start/stop the `lpsched` daemon which controls all routing of print requests to printers according to priority.
- Jobs stopped in the middle when the spooler is shutdown will restart *from the beginning* when the spooler restarts. The LP spooler only knows if a job is in the queue or not. It has no concept of how much has been printed.
- The enable and disable options in SAM always perform `accept/enable` and "reject/disable" pairs. There is no way in SAM to keep just one printer from printing without making it reject requests also.

Teaching Tips

- You should go through some of the manual commands if the students are interested in them. If they think SAM will handle their needs (in most cases it will) just briefly mention the manual commands.
- `cancel` is the only way, however, to cancel print requests.
- `lpmove` is the most common way to move print requests between queues.
- If students ask if there is any way for a user to move requests the answer is yes. Any user can modify *their own* requests' information using `lpalt`.

Transition

What if our spooler is acting funny and things are not working correctly?

6-6. SLIDE: What If Something is Wrong?

What If Something is Wrong?

- One or more printers will not print.

  Is the spooling system active? (Is the scheduler is running?)
  Are the printers disabled?
  Are the printers connected properly and On-Line?

- Users are getting error messages from `lp`.

  Is the printer accepting requests?
  Is there a default system printer?

- Requests go through the spooler but nothing prints.

  Is the printer connected?
  Is the printer configured for the proper interface type?
  Is the device file correct?

If all else fails, remove then re-add the printer to the spooler.

50722 6-6 63

Student Notes

The slide lists some common spooler and printer problems and things to check if you are having trouble. The `lpstat -t` command can be used to answer many of the questions regarding the spooler status.

Reference

The *Installing Peripherals* manual has a great deal of configuration and connection information for each type of printer.

The *Error Message Catalog* has a chapter on error messages produced by the LP Spooler system. It also lists remedies for many of the problems.

6-6. SLIDE: What If Something is Wrong? Instructor Notes

Purpose

These are some very basic but common problems problems that administrators have with the spooling system.

Key Points

- There is usually no way for the system to tell if a serial device is accessible on the system. If a serial printer is not connected or is not turned on, the spooler will continue to send requests to that printer and they will be lost.
- Printers like LaserJets have both serial and parallel ports. If the printer is not properly configured to use the correct interface, it will not print and the request will be lost.

Teaching Tips

- You may want to ask the students for the command sequences or SAM procedures that will answer the troubleshooting questions on the slide. This could also server as good review of the section.

Transition

You may sometimes need to remove a printer from your system.

6-7. SLIDE: Removing a Printer

Removing a Printer

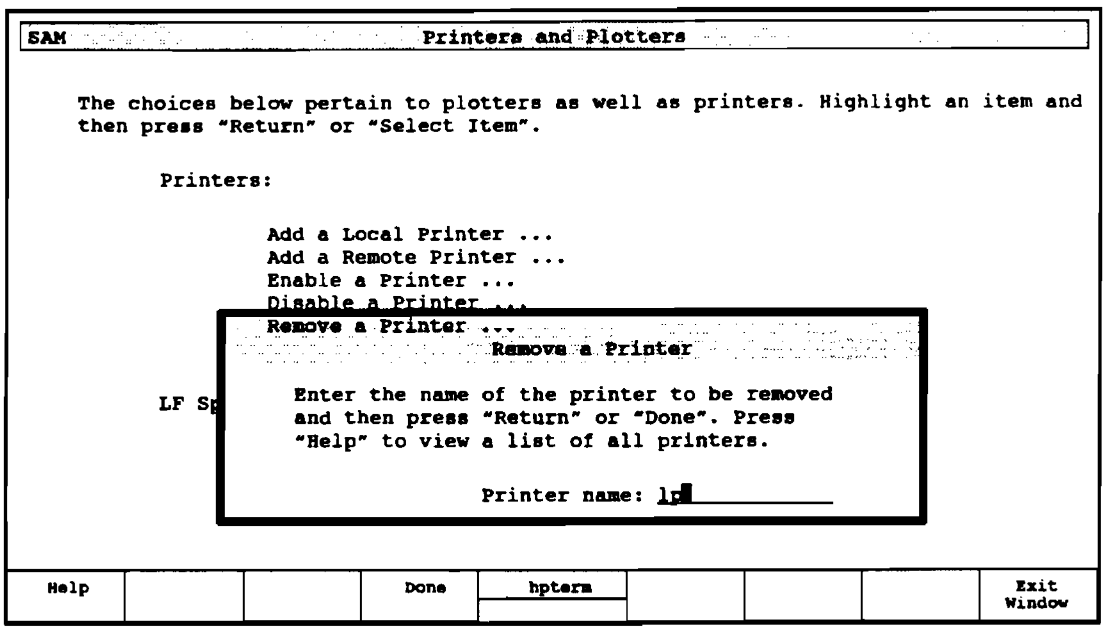

50722 6-7 64

Student Notes

To remove a printer with SAM, follow this procedure:

```
# sam (Return)

    Peripheral Devices ->

  Printers and Plotters ->

     Remove a Printer ...
```

SAM will prompt for the name of the printer to be removed as seen on the slide.

When you remove a printer with SAM all requests in that printers queue will be canceled. If you want to save these requests you must move them to another printer using `lpmove` before you run SAM. Remember that only `root` can run `lpmove`.

All information about the printer will be removed from the spooler. The `lpstat -t` command will no longer contain any information regarding the removed printer. The device file that was associated with the printer will *not* be removed from the system.

6-7. SLIDE: Removing a Printer

Instructor Notes

Purpose

Sometimes the spooler gets very confused and the only way to fix a printer is to remove it. Sometimes you just want to remove a printer that is no longer connected to your system.

Key Points

- Requests for the printer will be canceled.
- The device file is not removed.

6-8. WORKSESSION: Review Questions

1. What functions does the spooling system provide, and why are they required?

2. How would you cancel your own print request?

3. How would you cancel a print request owned by someone else?

4. How would the owner of a canceled request know that you had canceled it?

5. How can you cancel all requests for a printer using one command line?

6. What is the model/interface program used for?

7. If you have shut down the spooling system using SAM do any printers continue to print?

8. If you have stopped the spooling system using SAM can you still submit requests using `lp`?

9. How can you tell other users that a printer is "broken"?

10. How can you move print requests from the "printer1" queue to the "printer2" queue?

11. How can you move one print request (request-id "printer2-10") from "printer2" back to "printer1"?

12. What happens to requests in the queue of a printer that is being removed?

6-8. WORKSESSION: Review Questions

Instructor Notes

1. What functions does the spooling system provide, and why are they required?

Answer:

The spooling system manages print requests. They are required for optimizing the usage of limited system resources.

2. How would you cancel your own print request?

Answer:

```
# cancel request_id_number
```

3. How would you cancel a print request owned by someone else?

Answer:

You would have to find the request-id number then execute cancel as normal: `cancel` *request_id_number*

4. How would the owner of a canceled request know that you had canceled it?

Answer:

By a mail message, automatically generated by `cancel`.

5. How can you cancel all requests for a printer using one command line?

Answer:

`cancel` *printer_name* `-e`

6. What is the model/interface program used for?

Answer:

This is what sets up the printer to print your requests according to the `lp` command line options. There is a model/interface program for each type of supported printer in the `/usr/spool/lp/model` directory.

7. If you have shut down the spooling system using SAM do any printers continue to print?

Answer:

No. There may be information in the printer's buffer that will continue to print but no more information is sent to the printer from the spooler.

An important side effect is that any jobs that were printing when the spooler was stopped will start again *FROM THE BEGINNING* when the spooling system is restarted.

8. If you have stopped the spooling system using SAM can you still submit requests using `lp`?

Answer:

Yes. Requests can still go into the queue but they will not print until the spooling system is restarted.

9. How can you tell other users that a printer is "broken"?

Answer:

By `disable -r"Printer is broken"` *printer*. You could also disable the printer using SAM but you could not supply the specific reason the printer has been disabled.

10. How can you move print requests from the "printer1" queue to the "printer2" queue?

Answer:

`lpmove printer1 printer2`

11. How can you move one print request (request-id "printer2-10") from "printer2" back to "printer1"?

Answer:

`lpmove printer2-10 printer1`

12. What happens to requests in the queue of a printer that is being removed?

Answer:

The requests are all canceled.

6-9. LAB: Exercises

Directions

Work with your lab partner(s) to complete the following exercises.

If you are on a Series 800 you will have to coordinate your activities with the other users. Ask your instructor how this can best be accomplished.

1. Physically connect a printer to your system (if it is not already connected.)

2. Add the printer to your system using SAM.

3. Exit SAM. Submit a request to your new printer using the `lp` command. Did it print? If not, see if you can figure out why.

4. Check the status of your printer. Is your request listed?

5. Disable the printer. Now send some output to the printer. Now can you see your queued request?

6. Submit more requests to the disabled printer. Now cancel one of the requests from the queue. Check that the requests are gone.

6-9. LAB: Exercises

Instructor Notes

Lab Tips

If you do not have enough printers for each system you can pass a single printer around or if your systems are networked add the printer to one system as a group demo and then let the other system use SAM to add a remote printer as a bonus topic!

Alphanumeric terminals make pretty good 'dumb' printers also.

On the Series 800 you will have to coordinate the student's activities somehow. If you have the pseudo superuser environment each group of students can control their own spooling system independently at the same time. If you are not using the pseudo superuser environment, you will have to do the exercises in small groups one at a time.

1. Physically connect a printer to your system (if it is not already connected.)

Answer:

Connect the printer. Make sure to note the interface type and check if the printer is properly configured for that type of interface.

2. Add the printer to your system using SAM.

Answer:

See "Adding a Printer With SAM" in this chapter for details.

3. Exit SAM. Submit a request to your new printer using the `lp` command. Did it print? If not, see if you can figure out why.

Answer:

```
lp /etc/passwd
```

4. Check the status of your printer. Is your request listed?

Answer:

```
# lpstat -t
```

Your request will probably not be listed because it already printed.

5. Disable the printer. Now send some output to the printer. Now can you see your queued request?

Answer:

```
# disable -r"Printer down for testing" your_printers_name
# lpstat -t
```

The printer will be marked as "disabled" and your request should now be listed in the queue.

6. Submit more requests to the disabled printer. Now cancel one of the requests from the queue. Check that the requests are gone.

Answer:

```
# cancel request-id
request-id canceled.
```

`lpstat -t` will show you what you that the request were indeed canceled.

Module 7 — Terminals and Modems

Objectives

Upon completion of this module, you will be able to do the following:

- Add a terminal or modem to the system using SAM
- Describe what SAM does to add a terminal
- Use the `tset` command to initialize terminal characteristics.

About This Module

This module describes how to add a terminal or modem to a computer system using SAM. It also explains how to initialize terminal characteristics using the `tset` command to ensure that your applications will be able to interact with your terminal correctly.

Module 7 — Terminals and Modems

Overview of Module 7

This module describes how to configure a terminal or modem on the system. It uses SAM, of course, to add the terminal but there is also a brief discussion of how the terminal actually works with `init` and `getty` this is done to understand how to get terminals working in run-levels other than 2 and to better be able to troubleshoot terminal problems.

The `tset` section was purposely kept simple. There are many more things `tset` can do but they are probably beyond the needs of this audience.

7-1. SLIDE: Terminals and Modems

Terminals and Modems

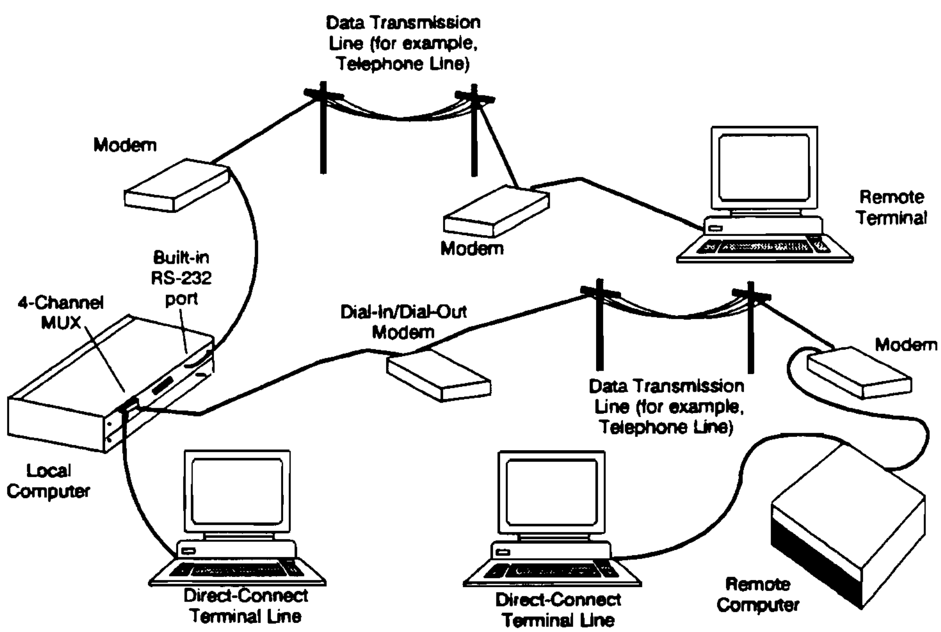

50722 7-1 65

Student Notes

Adding a terminal to your system allows you to log into your system from somewhere other than the system console. A modem may be used by a remote terminal to access your local computer. It also allows you to access remote computers via the telephone lines to share information. Users of a remote computer can also access your local computer with a modem. Terminals and modems are serial (RS-232) devices.

Each terminal or modem you connect to your computer takes up one serial port. This may either be an internal serial interface or a serial port on a MUX. Modems also require that the serial port support "modem-control". All current internal serial ports and most MUX ports can support a modem. Check with the *Installing Peripherals* for supported interfaces.

Note that one of the terminals connected to your system may be used as the system console. This will be the case on almost all Series 800s. On workstations the graphics monitor is more frequently used as the system console.

Reference

Installing Peripherals has information on serial interfaces and cables to connect different types of devices.

The owner's guides for the various models of SPUs will also have some of this information.

7-1. SLIDE: Terminals and Modems

Instructor Notes

Purpose

To describe the various ways in which terminals and modems may be used and to review how to connect terminals and modems to I/O interfaces.

Key Points

- A terminal allows users to access the system from somewhere other than the system console.
- On a Series 3/400 with a 4-port DIO MUX only port 0 supports a modem.
- On a Series 800 a terminal or modem may be connected to a MUX port. On the Nova (8X7) systems the system console is port 0 of the first ADP and the modem should be configured on port 7 to allow remote console access. See the 8X7 owner's guide for more info.

Teaching Tips

- You may want to poll the class to see how many students are using terminals and modems. On workstations, modems are more common than terminals.
- If you have students using DTCs (Data Terminal Controllers) you may want to mention that terminals connected to a DTC are not configured like a terminal connected to a local serial port. If a terminal is connected to a DTC and accessing an HP-UX system it is working as a `telnet` LAN session. You cannot use SAM to configure DTC based terminals and the only information in this chapter that applies is the terminal type setup. DTC setup requires either the *Entry Level DTC Manager* software for HP-UX or the *OpenView DTC Manager* software for a PC. The ARPA network services must also be installed on HP-UX for a DTC to work.

Preparation Notes

The configuration guides for the different SPUs have some of the most specific information regarding serial interfaces and cables used to connect the various devices.

Transition

Let's take a look at how to add a terminal or modem with SAM.

7-2. SLIDE: Adding a Terminal or Modem With SAM

Adding a Terminal or Modem With SAM

```
SAM                        Add a Terminal or Modem

    Fill in or modify the desired fields and then press "Perform Task".

                     Usage (mark one with an "x")  x  terminal
                                                   _  modem

                     Hardware Path . . . . . . . . 8

                     Port number . . . . . . . . . __

                     Speed (baud)  . . . . . . . . 9600

 Help | Main | Shell | Perform | hpterm |  |  |  | Exit
      | Menu |       | Task    |        |  |  |  | Task
```

50722 7-2 66

Student Notes

Before you use SAM to add your terminal, you must connect the terminal to your system. Make a note of the select code (Series 300/400), port number (Series 700), or hardware address and port number (Series 800) of your new terminal. Once the terminal is physically connected to the system, you can use SAM to make HP-UX recognize the terminal.

To get the menu shown in this slide select the following:

```
# sam (Return)

    Peripheral Devices ->

  Add a Terminal or Modem ...
```

Complete the menu with appropriate information. The following section describes in detail each field in the menu. Note that the second field varies slightly based on machine type.

- "Usage" — Indicate whether you will be adding a terminal or modem by placing an x next to the appropriate device type.

  If you select modem, SAM will ask you:

  ```
  Do you want the device for calling out? (y or n)
  Is this a CCITT modem? (y or n)
  ```

  If you want to use the modem to call out to other computers (as well as to receive calls) you should enter y in response to the first question; otherwise, if you want to use the modem only to receive calls from other computers, enter n.

  If your modem is located outside the United States or Canada you must operate the modem according to the CCITT interface standard. Enter n in response to the second question if your modem is located in the U.S. or Canada.

- Interface location. Only one of these three fields will appear depending on which machine type you are using:
  - *Select Code* — If you are executing SAM on a Series 300 or 400 you will be prompted for the select code that corresponds with the I/O interface used by your device. Using the `Help` softkey when your cursor is on this item will display a list of serial interface select codes.
  - *Card Slot Number* — If you are executing SAM on a Series 700 you will be prompted for the card number of the interface that the terminal or modem is connected to. Recall that the card slot number for any of the interfaces on the core I/O board will be 0 (zero).
  - *Hardware Path* — If you are adding this terminal to a Series 800 SAM will list all of your configured MUX devices before you enter this screen (see figure 7-1). When you select one of the listed MUXs, SAM will fill in this field for you. Reference the *System Hardware & Peripherals* module or *Installing Peripherals* for a review of how to read the hardware addresses.

- *Port Number* — This is the port that your terminal or modem is connected to. If you are using a Series 300 or 400 internal serial interface this number will be 0 (zero). On a Series 700 the port is either A or B for two internal serial ports. If your terminal is on a MUX this is the port number of the MUX.

  Press the `Help` softkey when your cursor is on this item and SAM will display a help screen containing valid port numbers (or letters).

- *Speed (baud)* — This is the rate at which data is transmitted to and from your terminal or modem. The unit used to denote speed of data transmission is called baud (bits per second). Press the `Help` softkey when your cursor is on this item and SAM will display a help screen that lists the valid baud rates that you may use. Terminals that are directly connected to the system are typically set to 9600 baud. Refer to the device's installation guide for help in selecting a baud rate.

Press `Perform Task` when you have provided all the information needed. SAM will create the device file if necessary and add the device to the /etc/inittab file. We saw /etc/inittab briefly in the *System Startup and Shutdown* chapter.

As soon as SAM completes the task you should see a login prompt on the new terminal (if you added a modem it will now be active).

Currently Configured MUX Cards

The following list shows ALL of the MUX cards SAM can find on your system. Move the cursor to the card where you wish to add the device and press "Select Item". If the MUX card you wish to use is not on the list, press "Device Missing".

—HARDWARE PATH—			
Card Slot	Bus Address	Driver	Description
8	–	mux2	HP–PB MUX Card
56	–	mux2	HP–PB MUX Card

Help			Select Item	hpterm		Device Missing	Dev File Info	Exit Window

Series 800 MUX Selection Screen

7-2. SLIDE: Adding a Terminal or Modem With SAM

Instructor Notes

Purpose

To show the students how to add a terminal or modem to their system using SAM.

Key Points

- Before using SAM to add a terminal or modem, you must know which I/O interface your device is connected to, and the select code or logical unit number of the interface. You also need to know which RS-232 port your device is connected to.
- `Help` is very useful on this screen.
- On a Series 3/400 8 port mux you actually get 2 sets of 4 ports: Select code 28 (by default) has the first four ports numbered 0-3 and select code 29 has the second four ports also labeled 0-3.
- CCITT is the Consultative Committee for International Telegraphy and Telephony)

Transition

The next slide describes what SAM does to add a terminal or modem to the system.

7-3. SLIDE: What Makes a Terminal Work?

What Makes a Terminal Work?

- The terminal's entry in `/etc/inittab` causes a `getty` to run:

  `a0:2:respawn:/etc/getty -h tty00 9600`

- The `getty` displays the `/etc/issue` file and displays the `login:` prompt.
- If the `getty` is not running, the terminal cannot be used to log in.

50722 7-3 67

Student Notes

When you add a terminal or modem, SAM creates a `getty` entry for the device in the `/etc/inittab` file.

Then SAM has `/etc/init` read the `/etc/inittab` file in order to activate the new `getty` process. Each terminal or modem connected to the system must have a `getty` process if users are to login there. You can identify which entry corresponds to the different terminals by the `ttyXX` name in the entry.

The `getty` entries in `/etc/inittab` are normally activated when the system is booted.

Note

HP-VUE Users: Recall that your system is not set up to run HP-VUE by default. If you modified `/etc/inittab` to automatically run HP-VUE (by putting a 3 or 4 in the `initdefault` entry) then you will have to add the 3 (and/or 4) to the second field of `inittab` entry for the new terminal to make it work:

```
a0:234:respawn:/etc/getty -h tty0p0 9600
```

Troubleshooting Tips

If the terminal is not working you can execute the ps command to see if the `getty` is running on the device.

```
# ps -ef

  UID   PID  PPID  C    STIME TTY      TIME COMMAND
     .
     .
 root     1     0  0  May  6  ?        0:20 init
 root   226     1  0  Jun  6  tty0p0   0:00 /etc/getty -h tty0p0 9600
 root   227     1  0  Jun  6  ttyd0p1  0:00 /etc/getty -h ttyd0p1 1200
 root 16572     1  0  Jun 18  tty0p2   0:00 /etc/getty -h tty0p2 9600
 root 16573     1  0  Jun 18  tty0p3   0:00 /etc/getty -h tty0p3 9600
     .
     .
```

If it is not running, you can try to start it using the command `init q`. This will force `init` to reread the `inittab` file and try the `getty` again. If the `getty` is still not running try to reboot your system.

If the `getty` is running but the login prompt is not being displayed, check that the terminal is properly connected (with the correct cable) and in **Remote Mode** (see *Installing Peripherals* on how to set your HP terminals' modes). If the terminal is set up properly, try to kill the `getty` process using the `kill` command:

```
# kill PID_of_getty
```

If the terminal still does not work you may have more serious trouble. Contact HP.

Removing the Terminal

If for some reason you no longer want the terminal or modem on your system, simply delete its line from `/etc/inittab` or change the word `respawn` to `off` then execute `init q`.

7-3. SLIDE: What Makes a Terminal Work? Instructor Notes

Purpose

To describe a little about what makes a terminal work and how to do some basic system troubleshooting.

Key Points

- SAM only adds terminals in run-level 2 (a 2 in the second field of the `inittab` entry). Normally HP-VUE is started by putting the system into run-level 3 (run-level 4 if the console is not a graphics display and VUE runs on X terminals). If you also want terminals or modems active in run-level 3 (or 4) you must modify the `inittab` entry.

Teaching Tips

- The troubleshooting tips are for some of the more basic problems like the getty process getting "hung" or changing speeds accidently. We won't even try to get into hardware problems.
- You might have students do a `ps -ef` to see a getty is running, `kill` it then do a `ps -ef` again. They should see a different process-id (PID) on that device's getty. When you kill a getty process it will restart again because of the `respawn` entry in `/etc/inittab`.

HP Terminal Modes

The *Installing Peripherals* manual goes into detail about setting HP terminal modes. A brief overview is given here.

To set terminal modes:

1. Press the (User/System) key at the top of your keyboard to display the terminal's "system" softkeys.
2. Press `Modes` to display the different modes.
3. To properly operate with HP-UX only the `Remote Mode` softkey should have a * in it. Pressing any mode softkey will toggle the mode on or off.

`Auto LF` and `Memory Lock` often get activated when the user displays a non-text file to their screen. Strange things will happen if these are on (have a * in them).

Pressing (User/System) then `Config Keys` will bring you to the terminal configuration options. `Datacom Config` is how you set speed and parity. `Terminal Config` and `Global Config` also have useful configuration options. Try them!

A Problem with Speed Settings

If a getty process is running on a terminal and it sees a Break on the line, it will change its speed according to the definition in the /etc/gettydefs file. Unfortunately, turning off a terminal is the same as pressing Break. So if you shut off a terminal while the getty is running, the getty will change speeds. When you turn the terminal back on, you will not see a login prompt because the getty is operating at a speed different from the terminal. At this point you can either kill the getty process so it will restart at the proper speed or you can press Break several times slowly to bring the getty back to the correct speed.

To avoid this type of problem completely you must edit the /etc/gettydefs speed definition so it will not change speeds. The 9600 entry in gettydefs normally cycles to 300 when getty sees the Break. To make it cycle back to 9600 instead, change the entry so it looks like this:

```
9600    # B9600 HUPCL IGNPAR PARENB ICRNL IXON OPOST ONLCR CS7 CREAD
                      ISIG ICANON ECHO ECHOK PARENB ISTRIP IXANY TAB3
        # B9600       SANE CS7 PARENB ISTRIP IXANY TAB3 HUPCL
        #login: #9600
```

Note that the only thing that changed is that the last 9600 used to be 300.

See gettydefs(4) and getty(1M) for more information.

Transition

We have already discussed setting up a user's shell environment. We said that one very important part of the environment is the terminal type.

7-4. SLIDE: Setting the Terminal Type

Setting the Terminal Type

Your terminal type can be set:

- using `tset` in your login script to prompt you at login:

  TERM = (hp)

- using `tset` in your login script to automatically set the type:

  No prompt

- anytime by manually setting the TERM environment variable:

  TERM=*type*
  `tset`

50722 7-4 68

Student Notes

In order to communicate with your terminal correctly, the HP-UX system must know the type of terminal or graphics display you are using when you log on.

The TERM environment variable represents the terminal type. If any command needs to control your terminal or know its size (like `vi` will do) you need have the TERM environment variable set correctly.

TERM is normally set using a `tset` command in every users' login file (.profile or .login). The default user login files will prompt the user for a terminal type when they login:

```
TERM = (hp)
```

Pressing Return, sets the TERM environment variable to hp, the basic terminal characteristics for any HP terminal. If you are on a special HP terminal (like a graphics display that is not running a window system) or if you are not on an HP terminal, you simply enter your correct terminal type then press (Return).

Valid terminal type names can be found under the `/usr/lib/terminfo/`*X* directories. Where *X* is the first character of the terminal name.

In general, if you are on an alphanumeric terminal, you can look in the directories under `/usr/lib/terminfo` for the terminal's model number to find the correct terminal type. For example, the "HP 700/92" terminal could use the terminal type `hp`, `70092`, `70092A` or `70092a`.

Setting your terminal type automatically at login

The `TERM = (hp)` prompt that you see after logging onto your terminal is generated by the following command in your login script (`.login` or `.profile`):

```
eval ' tset -s -Q -m ':?hp' '
```

If your users always use the same type of terminal, you can replace the ?hp with the valid terminal type. For example:

```
eval ' tset -s -Q -m ':70092' '
```

will set the user's `TERM` variable to 70092 whenever s/he logs in without asking any questions.

Reference

There are many other tricks you can use with `tset` to set your terminal type differently based on what terminal you used to log into the system. See the `tset` entry in the *HP-UX Reference Manual* for more details.

Setting the TERM Type Manually

If for some reason you did not get the correct `TERM` type set when you logged in, you can set it manually using the following commands:

```
$ TERM=type
$ tset
```

Where *type* is, of course, the correct terminal type. Executing `tset` resets your terminal characteristics based on the value set in `TERM`.

7-4. SLIDE: Setting the Terminal Type — Instructor Notes

Purpose

To help students set the TERM variable correctly for new user accounts that they have created.

Key Points

- This may seem like a lot of work for a simple variable setting but it pales in comparison to the trouble that can be had if the users' applications do not work.
- We purposely left out the details of `/etc/ttytype` because it can get very confusing and is seldom very useful to the type of administrator in the audience. If you feel the class could handle customizing `tset` and `ttytype` for things like dialing or network access, do it as an addendum to this section. Be careful not to confuse people that do not really care to see these advanced uses of `tset`.

Teaching Tips

- Most students should be able to hard code their types.
- If you have students that log into graphics displays in text mode (displays that are not running HP-VUE or X), you should mention the `TERM` types of those displays:

Table 7-1.

Machine Type	Resolution	Number of Text Columns and Lines	TERM Type
Series 300/400/700	1024x768	128x46	300h
Series 300/400/800	1280x1024	128x49	98550
Series 700	1280x1024	128x49	a1096

  There may be other type names based on graphics accelerator type. Your best bet is to search through `/usr/lib/terminfo` for something that matches the accelerator's part number.

- Now might be a good time to describe how to fix the superuser's .profile (`/.profile`) so this question is not displayed whenever root logs in:

```
Is your console one of the following: a 2392A, 2393A, 2397A or 700/92? [y/n]:
```

  To avoid this question being asked you must:

  1. Either enter the valid terminal types into `/etc/ttytype`:

```
70092 console
```

```
70092 tty0p1
```

Or

Change the `tset` line (near line number 240) to look like one in the student notes:

```
eval ` tset -s -Q -m ':70092' `
```

2. Then you must Change the `ask=yes` line (near line number 246) to:

```
ask="no"    # change to "no" if sure of value in /etc/ttytype
```

7-5. WORKSESSION: Review Questions

Directions

Complete the following questions.

1. What information does SAM need to know in order to add a terminal or modem?

2. What are some of the tasks that SAM performs to add a terminal or modem?

3. How do you set a terminal's type? Why is it necessary to tell the system the terminal's type?

7-5. WORKSESSION: Review Questions

Instructor Notes

1. What information does SAM need to know in order to add a terminal or modem?

Answer:

SAM needs to know the select code (if system is a Series 300 or 400), the card number (if the system is a series 700), or the logical unit number (if system is a Series 800), the port number, and the baud rate of the device you are adding.

2. What are some of the tasks that SAM performs to add a terminal or modem?

Answer:

SAM creates the necessary device file(s) if they do not already exist, adds a `getty` entry to the `/etc/inittab` file and SAM also has the system re-read the `/etc/inittab` file to activate the new `getty`.

The newly started `getty` process displays a `login:` prompt on the terminal and waits for a user to log in.

3. How do you set a terminal's type? Why is it necessary to tell the system the terminal's type?

Answer:

The terminal type is set each time a user logs onto a terminal, either automatically (via the `tset` command) or manually (using the `TERM=`*type* command).

The terminal type tells the HP-UX system about the terminal's display characteristics (e.g., size and other special characteristics like what escape sequence is used to clear the screen).

7-6. LAB: Exercises

Directions

You should have a terminal with your classroom system so you can practice adding and troubleshooting a terminal connection.

Once again you must coordinate your activities with your system partners.

1. Determine where the new terminal will be physically connected. What is the terminal's port number? If you have a Series 300 or 400 system, what is the select code of the I/O interface that the terminal is connected to? If you have a Series 700 what is the card number and the port letter? If you have a Series 800 system, what is the logical unit number of the I/O interface and port number that the terminal is connected to? What baud rate (speed) should the terminal use?

2. Physically connect the terminal then add the new terminal to your system with SAM and test it by logging on. Does it work?

3. Look at the `/etc/inittab` file and determine which `getty` entry relates to your new terminal port.

4. What is your new terminal's type?

5. Edit a user's `.profile` or `.login` script so that the system will automatically set the TERM type to the appropriate value for the terminal you are using.

Test your changes by logging in as the user and executing the `env` command. What the the value of TERM?

7-6. LAB: Exercises

Instructor Notes

Have your Series 300/400/700 students configure their systems to recognize a new terminal (using SAM) and then test out their configuration with a terminal. If you don't have several spare terminals, have the Series 300 students take turns testing their configuration with one spare terminal.

You will have to coordinate your Series 800 students. The pseudo superuser environment does not currently allow students to add terminals.

1. Determine where the new terminal will be physically connected. What is the terminal's port number? If you have a Series 300 or 400 system, what is the select code of the I/O interface that the terminal is connected to? If you have a Series 700 what is the card number and the port letter? If you have a Series 800 system, what is the logical unit number of the I/O interface and port number that the terminal is connected to? What baud rate (speed) should the terminal use?

Answer:

Answers will vary with machine type. Ask your instructor if you are not sure of any of the required values.

2. Physically connect the terminal then add the new terminal to your system with SAM and test it by logging on. Does it work?

Answer:

See "Adding a Terminal With SAM" for the details.

3. Look at the `/etc/inittab` file and determine which `getty` entry relates to your new terminal port.

Answer:

Execute the `tty` command (with no options) on the new terminal to determine its device file name. Search the `/etc/inittab` file for a `getty` entry that matches the device file name.

4. What is your new terminal's type?

Answer:

Look in the `/usr/lib/terminfo` database. The terminal type should closely match the product number of your new terminal.

5. Edit a user's `.profile` or `.login` script so that the system will automatically set the `TERM` type to the appropriate value for the terminal you are using.

Test your changes by logging in as the user and executing the `env` command. What the the value of `TERM`?

Answer:

[Edit your shell script to include a line like the following:]

```
eval ' tset -s -Q -m ':70092' '
```

Where 70092 would be whatever type of terminal you are using.

env should display TERM=70092

Module 8 — Disks and File Systems

Objectives

Upon completion of this module, you will be able to do the following:

- Add or remove a disk drive.
- Create a file system.
- Mount and unmount a file system.
- Define file systems to be automatically mounted.
- Determine free disk and file system space.
- Add dedicated and dynamic swap space.
- Add or remove a CD-ROM file system.
- Describe the disk layout of the different machine types.

About This Module

This module discusses everything you need to know to create a file system on disk. As system administrator, it is your responsibility to manage space in your file systems. You need to know how to monitor space and create new file systems when the current file systems are full. When you have several file systems, you need to know how to mount them and unmount them both manually and automatically.

Module 8 — Disks and File Systems

Overview of Module 8

If you are truly delivering this course to an operator audience that has nothing to do with disk configurations, then go through this section very quickly. Students should at least know how to view disk and file system information using `bdf` and SAM. They would do well to also know their disk layout and how to mount and unmount file systems. Even if they may never be adding a new disk to the system.

Preparation Notes

Work through the different SAM screens. There are many different ways to accomplish the same task. Know which screens do what. The notes try to bring out the most sensible way to configure disks and file systems. You can get into trouble if your are not sure what they all do. In general, if you are adding a new disk (no defined sections or file systems) use "*Add a Hard Disk*". Everything else can be done through the other screens.

The *HP-UX System Administration Tasks* manual and *How HP-UX Works: Concepts for the System Administrator* both have sections you should read to better understand file system structure and management and swapping and memory management.

The *PA-RISC Theory of Operations* course (H5081S) has very detailed information on these subjects. Far too much for the intended audience but good background information to help you as an instructor. You should be very familiar with HP-UX systems and be able to teach the 5-day system administration courses before tackling *Theory of Op*.

Warning **This is the most hardware dependent module (next to the hardware overview).**

The SAM screens can be quite different for different architectures. Try to de-emphasize the device names and hardware addresses and just look at the important information.

8-1. SLIDE: HP-UX Disks

HP-UX Disks

Disks are used for two things:

- File Systems — used to hold files and directories.
- Swap Space — used as an extension to system memory.

The tasks involved with management are the generally the same for all HP-UX systems but some details will differ due to the different disk layouts used on different architectures.

50722 8-1 69

Student Notes

The disks and file systems on your system are probably your most important resources. It is important that you understand the tasks involved with managing these devices. In this module we will discuss many tasks including:

- Adding new disks and creating new file systems.
- Mounting the file systems (making them available for use).
- Unmounting the file systems.
- Viewing file system and disk space information.
- Adding swap space to the system if necessary.

To accomplish these tasks we will be using SAM.

You will notice that when dealing with disks and file systems in SAM there is more than one way to accomplish the task. We will be discussing the most straight forward methods in SAM and we will try to give you general guidelines for when to use which method.

The disk layouts are different between the Series 800 systems, the Series 700, and the Series 300/400. These differences affect the screens in SAM enough to make it cumbersome to discuss all of these systems at once.

Therefore, we will cover the general concepts of file system and swap space management that apply to all systems first. We will then go into the architecture specific details and management differences. This will give us the basis we need to then proceed to the actual task of using and administering the disks and file systems.

8-1. SLIDE: HP-UX Disks

Instructor Notes

Purpose

To introduce the tasks at hand and the way they will be covered.

Key Points

- The disk sectioning scheme on the Series 800 makes quite a difference in some of the SAM screens. We will work around this and (hopefully) minimize the confusion.

Transition

File system management is a large part of the administrator's duties. To perform this duty competently, you must first understand what a file system is.

8-2. SLIDE: What is a File System?

What is a File System?

A way to organize a set of files and directories on a portion of a disk.

- A file system must be **mounted** to an empty directory.
- Access to the file system is through the mount point directory.

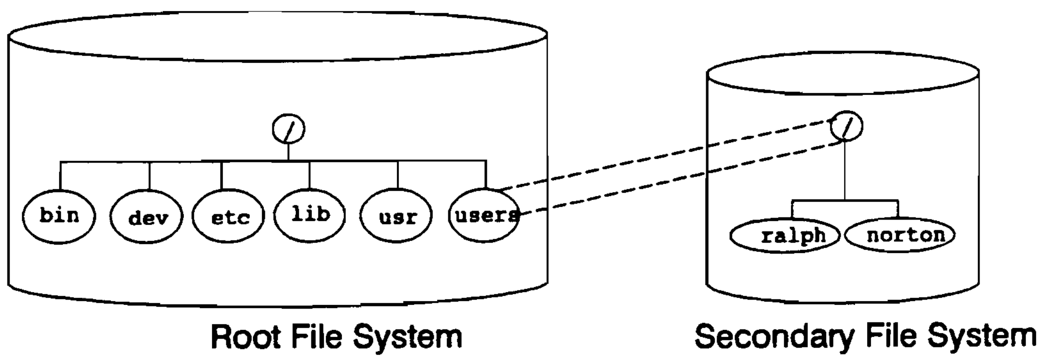

50722 8-2 70

Student Notes

Definition

A **file system** is a collection of files and directories that occupy a portion of a disk. **HFS**, **CDFS** and **NFS** are the three types of file systems HP-UX will support.

HFS The file system HP-UX uses on hard disks. It stands for the "High-Performance File System".

CDFS The file system HP-UX understands on a compact disks (called a CD or CD-ROM). The name means "Compact Disk File System".

NFS The "Network File System". This is used for sharing file systems with other computers on the network.

Every HP-UX computer has at least one HFS file system, the "root" file system. This file system contains at least the `/`, `/bin`, `/dev`, `/etc` and `/lib` directories. Other directories such as `/users` or `/usr` may be part of the root file system or they may be on a separate file system. They are commonly attached as separate file systems so they can be managed independently of the other file systems. For example, if your root file system is getting full, you may decide to move the `/usr` directory to a new file system with more space. The space on the root file system is now free for other uses. Secondary file systems might also contain files and directories for specific applications. For example, there is a CD-ROM that contains the entire HP-UX manual set and is accessed using the LaserROM/UX software. This file system can be attached and detached as needed without affecting the use or space of the other file systems.

HP-UX cannot use a file system unless it is **mounted**.

Definition

Mounting is the way a file system is incorporated into the overall directory tree. When a file system is mounted, it is attached to an empty directory on an already mounted file system. This directory is called the **mount point** directory.

Only the superuser can mount and unmount file systems.

Access to the file system is now as simple as using the mount point directory. The files and directories in the mount point directory are actually on the separate physical file system.

To disallow access to the separate file system we can unmount it and the mount point directory will be left empty.

The Series 300/400/700 systems can have only *one file system per disk*. The Series 800 may split a hard disk into several **disk sections**. *Each section* may contain a different file system. We will discuss the details of the different disk layouts later in this module. Throughout this module when we discuss a file system we are talking about a whole disk on a Series 300/400/700 system or a disk section on a Series 800.

Reference

The book *How HP-UX Works: Concepts for the System Administrator* has a detailed discussion of the HFS file system structure.

The *HP-UX System Administration Tasks* manual contains several sections on file system management. This would be a good reference for all topics in this module.

8-2. SLIDE: What is a File System? Instructor Notes

Purpose

To define a file system. We will be using the terms presented here throughout this and the next module. Make sure students understand them.

Key Points

- If the mount point is not empty, the files in it will be inaccessible until the file system is unmounted.
- There is one file system per disk on Series 300/400/700 systems. One per disk section (possibly many per disk) on a Series 800. The details of the layouts are covered soon.
- Only the superuser can mount and unmount file systems.

Teaching Tips

- Ask students how many of them have more than one disk on their systems. If they have more than one disk, they will most likely have more than one file system and the concepts presented regarding file systems will be very important to them.
- Point out and explain the example of putting `/usr` on a separate disk to better balance disk space needs.

Transition

The other major use of disks on HP-UX systems is for swap space.

8-3. SLIDE: What is Swap Space?

What is Swap Space?

Space on disk(s) used to extend the system's memory.

There are two types of swap space:

Device Swap — Fixed amount of space on a disk. The space is dedicated to swapping.

File System Swap — Space taken out of a file system for swapping. This space can grow and shrink dynamically.

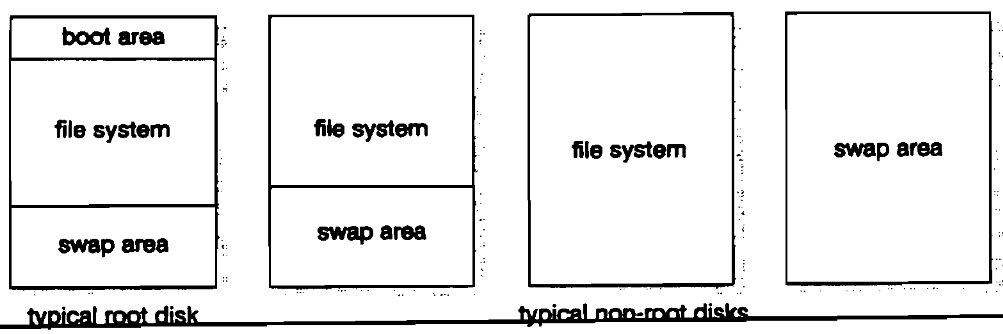

50722 8-3 71

Student Notes

We have mentioned swap space in relation to disks, disk sections and file systems. We now need to define it and better understand what it does for us so we can manage it properly.

Running programs need to have at least a portion of their code and data in physical memory in order to execute. Physical memory is a finite resource on a computer. This means that only so many processes can fit into physical memory at any one moment in time, even though many more processes may actually be ready to run. Because demand for physical memory may exceed actual supply, the concept of **swap** was introduced to computer systems. When demand for physical memory is high, entire processes or portions of processes are pushed out to the swap area on a disk (swapped out). When physical memory demand is low, portions of processes are brought back into memory from the swap area as needed.

Your system must be properly configured with enough physical memory and swap space to accommodate all of the requests for memory your applications will make. This information is normally supplied with the installation documents that come with your software. You should read these documents carefully and configure swap space appropriately before attempting to run large applications. If your system cannot handle all of the memory requests, the process requesting memory will die. This can cause very strange behavior in system tools like SAM.

Definition

There are some swap terms you should be aware of because they are used throughout the rest of this module and in the system documentation. They are **Device Swap**, **File System Swap**, **Primary Swap** and **Secondary Swap**. The definition of these terms is given below.

Device swap A disk or section of a disk that is used exclusively for swap. Device swap is always a fixed size and always dedicated to swapping. Another name for device swap is **dedicated swap**. There can be only one dedicated swap section per physical disk. (This is true for all HP-UX machines.)

File system swap File system swap, unlike the exclusively purposed device swap, is a file system that not only supports files and their data structures, but also has space available for swapping. The amount of space taken away from the file system for swapping can grow or shrink depending on the memory demands. That is why file system swap is often referred to as **dynamic swap**. File system swap is almost always slower than device swap because it has to deal with the file system. For this reason, it is not recommended that you rely heavily on file system swap in your configuration. If you absolutely need more swap space for your system to run applications correctly, you should add device swap.

Primary swap A special type of device swap which must be available at boot time. Primary swap is almost always located on the same disk as the root file system. Primary swap is configured when the system is first installed and cannot grow or shrink without rebuilding the system.

Secondary swap Device or file system swap that is used in addition to primary swap. It is typically located on disks other than the disk containing the root file system.

Primary swap is always activated at boot time. Secondary swap devices can be activated while the system is running as needed. We typically define all of the secondary swap space required to run our applications using SAM so they are activated at boot time also.

Once a secondary swap system is activated it cannot be deactivated. It will remain active until you reboot the system. If you are using file system swap to a mounted file system (like `/users`) you will not be able to unmount that file system without rebooting.

Reference

The *HP-UX System Administration Tasks* manual has a whole chapter on swap called "Managing Swap Space". In this section you will find more information on deciding how much swap you need and guidelines for configuring swap devices.

How HP-UX Works: Concepts for the System Administrator is another good source of information.

8-3. SLIDE: What is Swap Space? Instructor Notes

Purpose

This introduces the concept of swap, the different types and the fact that we can use SAM to configure our swap devices for automatic startup.

Key Points

- Device swap is always faster than file system swap.
- Once a swap device is active it cannot be deactivated.
- Primary swap must be available at boot time or the system will "`panic`" and crash.
- Secondary swap must be activated manually or configured to be activated at boot time.
- If the system appears to be doing very strange things like commands will not run once but if tried again they will work or if SAM is reporting strange internal errors you are probably running out of swap.

Teaching Tips

- Go through the terms. It is important that students understand these now.

Transition

Now that we have a basic idea of the parts of our disks we can take a look at the different ways the different architectures manage them.

8-4. SLIDE: Series 300/400 Disk Layout

Series 300/400 Disk Layout

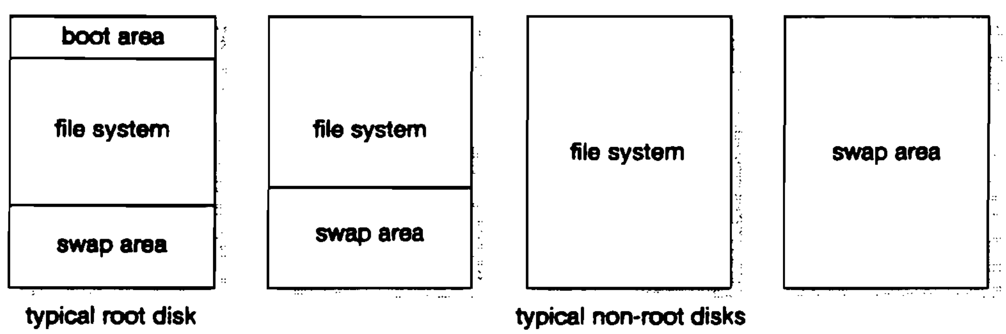

50722 8-4 72

Student Notes

On a Series 300/400, a hard disk can be comprised of a boot area, an HFS file system area, and a swap area. The **root disk** will contain all three of these areas.

Definition

The **root disk** or **system disk** is the disk that holds the root file system. It will be the disk the system initially boots and runs HP-UX from.

At the start of the root disk is a 8Kb area which identifies this as a bootable disk and contains the information necessary to boot from this disk.

Non-root disks will typically contain a single swap area, a single file system, or a combination of both.

Note

On the Series 300/400 computer there is only one file system area per physical disk. Any disk space not occupied by the file system (or boot area) may be used for swap space.

The root disk was built when HP-UX was first installed.

The configuration of non-root disks can be handled through SAM.

8-4. SLIDE: Series 300/400 Disk Layout — Instructor Notes

Purpose

To describe the Series 300/400 root and secondary disk layout possibilities.

Key Points

- One of the primary tasks of a system administrator is the creation and maintenance of file systems.
- The root disk layout is setup at operating system installation time and the proportion of the file system and swap space cannot be changed without reinstalling the software.
- Point out that the Series 300 and Series 400 disk layouts are the same.

Transition

The Series 700 disk layout is similar to the Series 300/400.

8-5. SLIDE: Series 700 Disk Layout

Series 700 Disk Layout

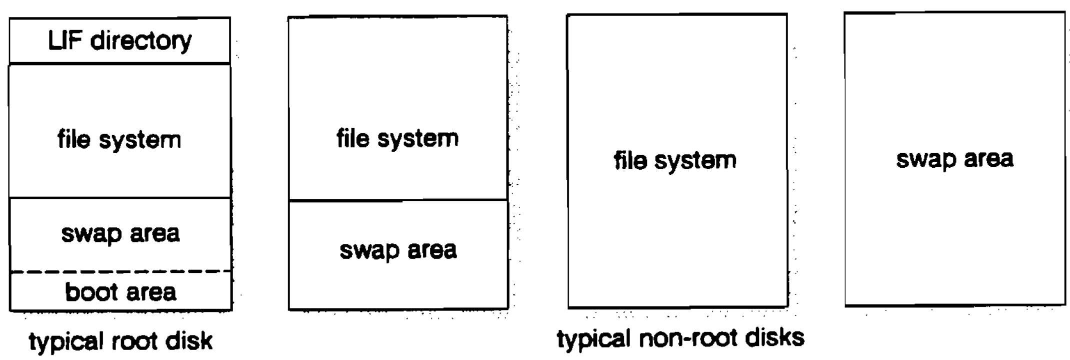

50722 8-5 73 ©1992 Hewlett-Packard Co.

Student Notes

On a Series 700, a hard disk can be comprised of a LIF directory, a file system area, a swap area, and a boot area. The **root disk** will hold all four of these areas.

Definition

The **root disk** or **system disk** is the disk that holds the root file system. It will be the disk the system initially boots and runs HP-UX from.

At the start of the root disk is a 8Kb area which is actually a small LIF (Logical Interchange Format) file system that is read by the boot ROM to find the boot area.

The boot area found on the last 2Mb of the root disk actually holds the bootstrap utility, the boot up code that will load in the HP-UX kernel and start it running.

Non-root disks will typically contain a single swap area, a single file system, or a combination of both.

Note

On the Series 700 computer there is only one file system area per physical disk. Any disk space not occupied by the file system and the 2Mb boot area may be used for swap space.

The root disk was built when HP-UX was first installed.

The configuration of non-root disks can be handled through SAM.

8-5. SLIDE: Series 700 Disk Layout

Instructor Notes

Purpose

To describe the Series 700 root and secondary disk layout possibilities.

Key Points

- One of the primary tasks of a system administrator is the creation and maintenance of file systems.
- The root disk layout is setup at operating system installation time and the proportion of the file system and swap space cannot be changed without reinstalling the software.

Transition

How do we get the file systems and swap space on a disk in the first place?

8-6. SLIDE: Adding a New Disk to Series 300/400/700 Systems

Adding a New Disk to Series 300/400/700 Systems

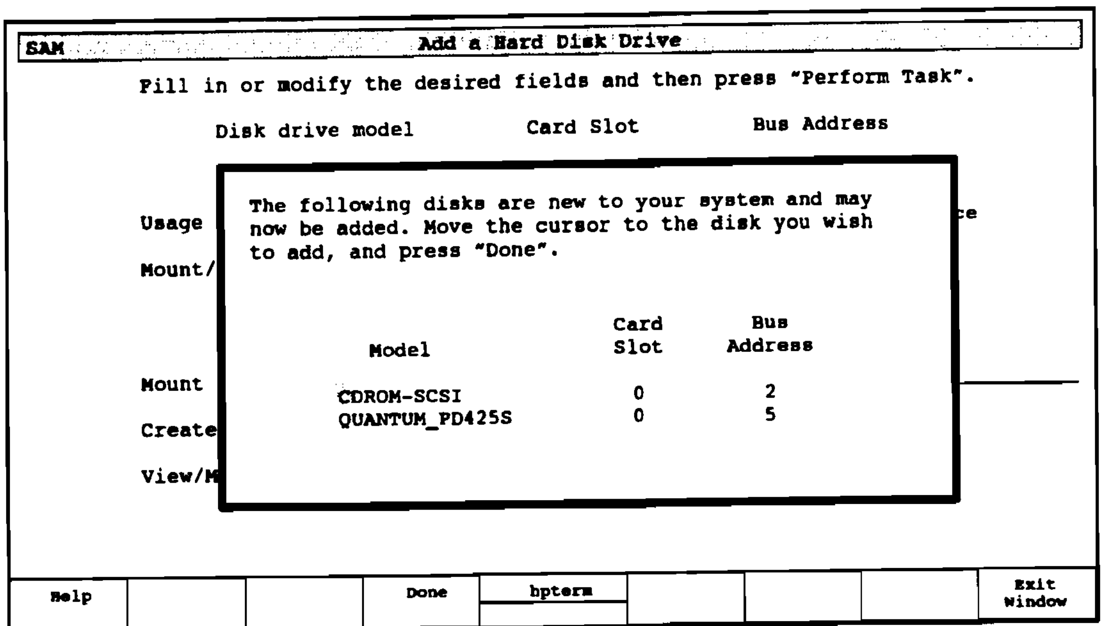

50722 8-6 74

Student Notes

As we already mentioned the root disk is completely laid out when the operating system software is first installed. But we also need to know how to assemble the secondary (non-root) disks.

This task is most easily done through SAM:

```
# sam (Return)

    Peripheral Devices ->

  Disk Drives ->

     Add a Hard Disk Drive ...
```

Once you make the "Add a Hard Disk Drive" selection, you can add a new disk to your configuration and optionally create a new file system on that disk. Creating a new file system on the disk will destroy any existing file system or data on the disk.

If the disk already contains a file system on it that you want to use (like a CD-ROM would), then you can either add it here and specify that you do not want a new file system created or you could add it through the "*Add a Local File System*" screen from SAM's `File Systems->` menu.

Before you can actually configure the new disk, SAM will display a popup window listing all the disks that are attached to your system but are not yet configured for use as a file system or swap space. The slide shows the screen used to select available devices.

You should choose the disk you want to add with the arrow keys and then press the `Select Item` softkey.

If your disk is not shown, you can select `Device Missing`. this softkey will cause SAM to present a number of items to check which will help to step you through the connection and configuration of your disk drive.

8-6. SLIDE: Adding a New Disk to Series 300/400/700 Systems

Instructor Notes

Purpose

To show the students that SAM can add new disks to our existing configuration. We also want to show that SAM can find all "unconfigured" devices for us. This leaves little guesswork to the task.

Key Points

- This screen is only to be used for adding a new disk to the system configuration. If a disk has already been defined on your system to contain swap space or a file system, it will not appear as an unconfigured device on this screen.
- The "*Add a Local File System*" screen is *only* used to add an existing file system into the configuration whereas "*Add a Hard Disk*" is intended to define a whole disk. It will build a file system on the disk and define the swap space if necessary. The "*Add a Local File System*" screen will not allow you to do any swap configurations.

Teaching Tip

- You should review the steps to physically add a hard disk here. This would be a good class review question.

  The steps for adding a disk are:

  1. Shut the system down and turn off the power.
  2. Choose the interface card to which the new disk will be connected.
  3. Choose a bus address for the new disk which does not match the bus address of any other device on the chosen interface card.
  4. Connect the new disk to the chosen interface card.
  5. Turn the new disk on.
  6. Turn the system on.
  7. Execute SAM to complete the process for adding the new disk.

Transition

We just selected the disk to add, now we need to define its usage to the system.

8-7. SLIDE: Defining a New Disk on the Series 300/400/700

Defining a New Disk on the Series 300/400/700

```
SAM                         Add a Hard Disk Drive

        Fill in or modify the desired fields and then press "Perform Task".

              Disk drive model          Card Slot          Bus Address

              QUANTUM PD425S            0                  5

        Usage (mark one or both with an "x") .  x file storage    x swap space

        Mount/enable when? (mark as desired) .  x now             x on boot

                If usage includes file storage, fill in the fields below.

        Mount directory . . . . . . . . . . . . /users

        Create a new file system? (y or n)  . . y

        View/Modify additional default file system options ? (y or n) n
```

Help	Main Menu	Shell	Perform Task	hpterm		Disk Info	File Sys Info	Exit Task

50722 8-7 75

Student Notes

Once you have selected the disk to add, you will see this screen.

The fields on the screen should be filled in as follows:

- *Usage* — This allows you to select if the disk is to be used for *file system*, *swap space* or both. You will be able to choose the file system size later on this screen.
- *Mount/enable when* — Marking *now* will mount the file system as soon as it is created. Marking *on boot* will define the file system to be automatically mounted at boot time.

The rest of the options apply only if you are allocating file system space on the disk.

- *Mount directory* — This is the mount point directory for the file system. The specified directory must be empty or SAM will warn you that the files in the directory will be inaccessible while the file system is mounted.
- *Create a new file system?* — If this disk is marked for *file storage* it needs a file system on it. Answering y to this question will display a window like this:

```
                       Create a New File System
Fill in or modify the desired fields and then press "Done".

Initialize disk? (y or n)  n

Disk space allocation:  50  M for swap,  353  M for file system


                    (use arrow keys to see more choices)
```

Initializing a hard disk formats the physical media and tests it so it can be used by HP-UX. You should initialize a hard disk if the disk is new or has never had HP-UX on it before.

To select the proportion between file system space and swap space, you can use the up and down arrow keys. When the combination that suits your needs appears, press Done .

- *View/Modify additional default file system options* — Answering n will accept the default values for these options.

Answering y to this item will bring up a popup screen that allows you to specify several configuration options.

```
Default File System Options

Fill in or modify the desired fields and then press "Done".

Write protection (mark one) . . . . . .  _ read only  x read/write

Set user ID execution allowed? (y or n)  y

Long file names? (y or n) . . . . . . .  n

Copy bootstrap program? (y or n). . . .  y
```

The options you can change are:

- *Write protection* — This allows you to select if the files and directories on the disk may be changed or not. Marking a file system as read-only indicates that you cannot edit, create, remove or rename files on that file system. This works well for CD-ROM file systems because they are inherently read-only. Most file systems will be marked read-write. Read-write is the default choice.
- *Set user ID execution allowed* — A set-user-ID program is a special type of program that turns you into another user while the program is running. This option allows you to specify whether users can run these special types of programs on the file system. If you are not sure about set-user-ID programs, mark the field with a y. Many of the system programs need this feature to operate correctly. The default choice is y.
- *Long file names* — The standard UNIX file name length is 14 characters. If you answer y to this question, this new file system will allow file names up to 255 characters long. If your applications can work on a long file name file system, you should answer y to this question. If you are not sure, answer n. If later you find that you would like to use long file names, you can convert the file system to support them. The default is to use long file names.
- *Copy bootstrap program* — Answer n to this question unless you are planning to build a bootable file system on this disk. The default answer is y.

Fill in the fields as described then press `Perform Task`. You will see several status messages on the screen. When SAM completes the task it will ask if you want to add another disk. If you answer n you will be brought back to the "*Disk and Swap Configuration*" screen.

The defined disk is now a part of the system configuration. If you re-enter this screen to add another hard disk, this disk will not show up as an available, unconfigured device.

If you decide for some reason that you want to remove the disk from the system, it can be done using the **`Remove a hard disk ...`** menu option.

Changing a File System's Size

If you want to change a disk's file system and swap proportions, you will have to backup any data on the file system then remove the disk from the system configuration. Later, you can add it again with the new parameters for file system and swap size. If you had a file system on the disk, it will be destroyed when you create the new file system. This is the reason for doing the backup before removing the disk. You simply have to restore the backup onto the new file system. This is the *only* way to change a file system's physical size.

8-7. SLIDE: Defining a New Disk on the Series 300/400/700

Instructor Notes

Purpose

To describe what needs to be entered to define the new disk.

Key Points

- SAM will create a file system and mount it if you choose.
- Once you define a disk in this fashion, you will not see it on this screen again unless it is explicitly removed. Using the `Remove a Hard Disk...` option.
- On the Series 700 SAM will also let you `Add a Floppy Disk ...` in much the same manor as described here. This is to support the optional internal 3.5 inch micro-floppy drives. The Series 300/400 cannot add floppy disks through SAM. Check the *Installing Peripherals* manual and the `/etc/disktab` file for details regarding adding floppies.

Teaching Tips

- Most of the fields are fairly obvious. You could describe most of them in very broad terms. The students can find the details in the notes.
- Many of these options are the same on a Series 800 but organized differently. You will see a very similar list in the Series 800 section. This is so you could deliver the course to an all Series 300/400/700 or all Series 800 audience and just skip whole pages that do not apply. If you have a mixed audience you only need to cover the details of the fields once.

Transition

The most complex disk layout scheme is on the Series 800.

8-8. SLIDE: Series 800 Disk Layout

Series 800 Disk Layout

Disks can hold:

- Boot area (2MB)
- File system
- Swap

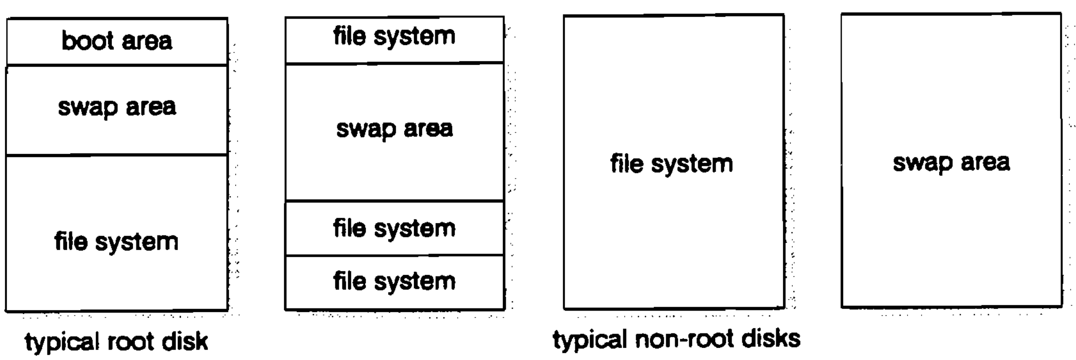

50722 8-8 76

Student Notes

On a Series 800, a hard disk can be comprised of a boot area, multiple file system areas, and a swap area. Each of these areas is known as a **disk section** or **disk partition**. Each disk section is treated as a separate device on the system.

The **root disk** typically contains the boot area, one file system section (sometimes more) and the primary swap section.

Definition

The **root disk** or **system disk** is the disk that holds the root file system. It will be the disk the system initially boots and runs HP-UX from.

At the start of the root disk is a 2MB area reserved for the boot programs. It is actually a simple file system called LIF (Logical Interchange Format). This LIF volume actually holds the boot up code that will load and execute the HP-UX kernel

Non-root disks can contain one or more file systems and possibly a single swap section. There should only be one swap section per physical disk.

The boot area, root file system and primary swap section are configured when you install the HP-UX operating system software. Other file system sections may also be configured at that time.

The configuration of non-root disks can be handled through SAM.

8-8. SLIDE: Series 800 Disk Layout

Instructor Notes

Purpose

To describe the various disk layouts for the Series 800.

Key Points

- The Series 800 supports disk sectioning. Each section is treated like a separate device with a separate device file.
- You technically do not need the boot area, root file system and primary swap on the same disk. Administration and kernel configuration is just made easier if it is. The installation process builds these three sections (boot, root and swap) on the disk being loaded.

Teaching Tip

- Do not get into disk sections too much here. The details are covered on the next slide.
- Make sure you mention that each section can be managed separately. This can allow us to have many file systems per disk.

Transition

Every disk on a Series 800 uses basically the same sectioning layout.

8-9. SLIDE: Series 800 Disk Sectioning

Series 800 Disk Sectioning

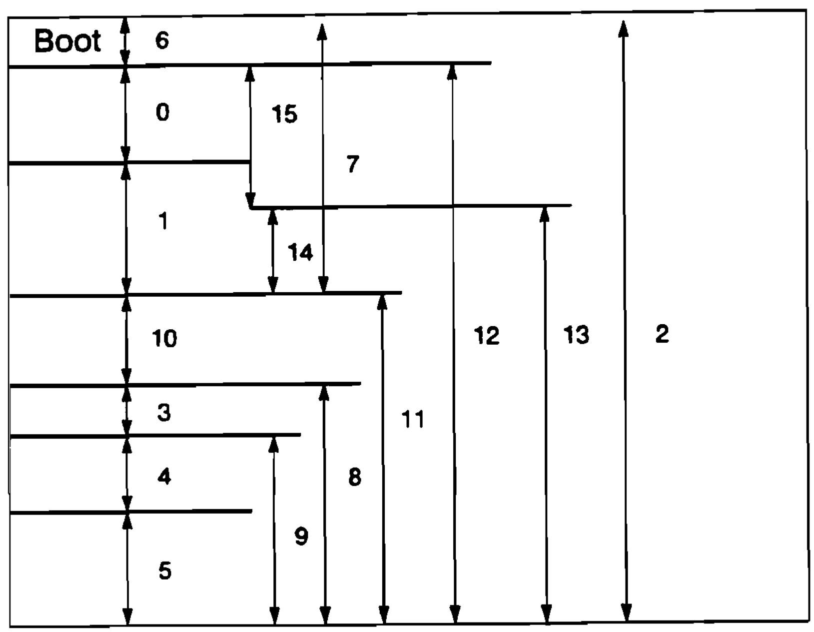

50722 8-9 77

Student Notes

This is a picture of the Series 800 disk sectioning scheme. It is almost identical for all disks used on these machines. The size of each section will vary by the total disk size. These are logical sections imposed strictly by the HP-UX software. Disk sectioning allows the operating system to treat each section as if it were a separate disk.

The advantages of employing a disk sectioning scheme generally outweigh the added complexity of managing such a scheme. Among the chief advantages is the ability to control the amount of disk space that can be used by a certain project or group of users. The administrator can select appropriate section sizes based on the application. Since a file system is contained in, and limited to the size of the disk section that it is created in, the file system cannot grow without bounds.

Notice on the slide that many of the disk sections overlap each other.

Warning

You may *NOT* simultaneously use disk sections that overlap. Doing so would eventually destroy data in the overlapping sections.

The HP-UX operating system will not prevent you from creating a file system in a section that encompasses a section already in use!

SAM, however, will impose this most important restriction.

The usage of the sections will vary from one installation to another. If a disk is bootable, however, the boot area will reside in section 6. This implies that the use of section 7 or section 2 on your boot disk is prohibited since section 6 is contained within sections 7 and 2. The root file system resides in section 13 by default. The primary swap space then is in section 15 (the only unused section remaining). This layout can be changed during the operating system installation procedure. The layout of the root disk sections (especially the root file system and primary swap sections) cannot be changed easily without reinstalling the system.

The number of sections that can be defined on a hard disk may vary with disk models. Disks smaller than about 350Mb will not have sections 10 or 11.

The `/etc/disktab` file defines the section sizes for the various supported disk models. A portion of the `/etc/disktab` file is shown here for a 1.3Gb hard disk:

```
C2474S:\
        :ty=winchester:ns#36:nd#19:nc#1935:rm#4002:\
        :s0#24280:b0#8192:f0#1024:\
        :s1#48560:b1#8192:f1#1024:\
        :s2#1323540:b2#8192:f2#1024:\
        :s3#29298:b3#8192:f3#1024:\
        :s4#107426:b4#8192:f4#1024:\
        :s5#981540:b5#8192:f5#1024:\
        :s6#1998:b6#8192:f6#1024:\
        :s7#75240:b7#8192:f7#1024:\
        :s8#1119024:b8#8192:f8#1024:\
        :s9#1089612:b9#8192:f9#1024:\
        :s10#129024:b10#8192:f10#1024:\
        :s11#1248300:b11#8192:f11#1024:\
        :s12#1321488:b12#8192:f12#1024:\
        :s13#1272924:b13#8192:f13#1024:\
        :s14#24280:b14#8192:f14#1024:\
        :s15#48560:b15#8192:f15#1024:
```

The first part of the entry is the disk model identifier (C2472S in this sample). The entries starting with `:s`*S*`:#` are the sizes of the different sections available on this disk given in Kbytes (*S* is the section number).

Note that section 2 is always the entire device and section 6 is always 2Mb regardless of disk size.

8-9. SLIDE: Series 800 Disk Sectioning — Instructor Notes

Purpose

Understanding sectioning is a crucial part of disk management. You cannot do anything to disks without knowing how the sectioning scheme works.

Key Points

- The layout of the root disk cannot easily be changed without reinstalling HP-UX.
- It is imperative that students understand that a file system is contained within disk section boundaries and that file systems must not be created in sections that overlap. Disaster will ensue if file systems are created in sections that overlap.
- The boot area resides on section 6. All boot disks have section 6. On disks that you do not boot from, section 6 could contain other information or be overlapped by using sections 7 or 2.
- Point out that section 2 always addresses the entire medium, regardless of any entry in `/etc/disktab`.
- With HP-UX (on the Series 800), section sizes are predetermined and cannot be changed. As a matter of fact, the range of cylinder addresses of each of the sections is fixed in the disk driver software. The only way to modify the given disk sectioning scheme is by recoding the driver, which is not supported.

Sectioning Note

The advantage of disk sections is that they allow the system administrator to manage disk space at a smaller granularity than the whole disk (as the workstations do). The system administrator can select different size disk sections for differing filesystem and swap space requirements. This allows more flexibility than other single filesystem implementations.

The implementation of HFS at the 8.0 release supports the notion of `disk quotas`. This is derived from the BDS 4.3 release, and allows the administrator to set specific limits for users on the amount of disk space they can consume. Quotas are implemented on a file system level. Disk sections can give the administrator more control over the quota definitions. We will not be covering quotas in this class.

Teaching Question

- What is the maximum number of file systems that could fit on a disk drive?

  (Answer: Seven. Five if this is the boot disk because a boot disk will also contain the boot area and primary swap by default.)

- If a file system completely filled section 1, could you make a file system in section 14?

  (Answer: Not without destroying the file system in section 1)

- Can a file system be larger than its section?

(Answer: No. File systems and swap areas are totally contained in their defined sections.)

Note The sectioning scheme picture presented at the beginning of the /etc/disktab file is *wrong*. The section information for each disk type is correct. Only the picture in the top of the file is wrong.

The corrected picture should look like this:

```
# This file contains the disktab entries for the current sectioning
# scheme
#
# Note that the layout of the disks may vary.
#
#   With disks SMALLER than ~ 350MB the following layout applies:
#
#       -----------------------------------------
#       6                         ^         ^
#       ----------------------    |   ---   |
#       0            15           | 7  ^    |
#                ------------     |    |    |
#       ---------            ^    |    |    |
#       1            14      |    v    |    |
#       ------------------   | --------|    | 2
#       3                ^   | 13      |    |
#       -----------      |   |         | 12 |
#       4      ^         | 8 |         |    |
#       ------ | 9       |   |         |    |
#       5      v         v   v         v    v
#       -----------------------------------------
#
#
#   With disks GREATER than ~ 350MB the following layout applies:
#
#       -----------------------------------------
#       6                         ^         ^
#       ----------------------    |   ---   |
#       0            15           | 7  ^    |
#                ------------     |    |    |
#       ---------            ^    |    |    |
#       1            14      |    v    |    |
#       ------------------   | --------|    | 2
#       10                   |    ^    |    |
#       ------------------   |    |    |    |
#       3                ^   | 13 |    |    |
#       -----------      |   |    | 11 | 12 |
```

```
#        4        ^        | 8  |      |     |     |
#        ------ | 9        |    |      |     |     |
#        5        v        v    v      v     v     v
#        -----------------------------------------
#
```

Transition

How can we build a non-root disk using this scheme?

8-10. SLIDE: Adding a New Disk to Series 800 Systems

Adding a New Disk to Series 800 Systems

Available Disk Drives

```
Here is a list of NEW disk drives SAM can find on your system. The list
also includes disks with at least one unused section. Move the cursor to
the desired disk drive and press "Select Item". If the disk drive is
not on the list, press "Device Missing".
```

—HARDWARE PATH—			
Card Slot	Bus Address	Interface	Description
4	1	hpib	HP-IB Disk
52	2.0	scsi	SCSI Disk
52	5.0	scsi	SCSI Disk

Help			Select Item	hpterm		Device Missing		Exit Window

50722 8-10 78

Student Notes

As we already mentioned the root disk is completely laid out when the operating system software is first installed. But we also need to know how to assemble the secondary (non-root) disks.

This task is most easily done through SAM:

```
# sam (Return)

    Peripheral Devices ->

        Disk Drives ->

     Add a Hard Disk Drive ...
```

Once you make the "Add a Hard Disk Drive" selection, you can add a new disk to your configuration and optionally create new file systems on a set of disk sections that you select.

If any disk section already contains file system(s) on it that you want to use (like a CD-ROM would), then you can either add it to the system configuration here and specify that you do not want a new file system created on that section, or you could add it through the *"Add a Local File System"* screen from SAM's `File Systems->` menu.

Before you can actually configure the new disk, SAM will display a popup window listing all the disks that are attached to your system that have sections that are not yet configured for use as a file system or swap section. The slide shows the screen used to select available devices.

You should choose the disk you want to add with the arrow keys and then press the `Select Item` softkey.

If your disk is not shown, you can select `Device Missing`. This softkey will cause SAM to present a number of items to check which will help to step you through the connection and configuration of your disk drive.

After pressing `Select Item` softkey, another popup will appear to choose which disk sections to use on the disk. If the disk is *new* to your configuration you need to select the sections using this screen:

```
            Choose Number of Disk Sections

Move the cursor to the number of sections you wish to
have on the disk and then press "Done".

SAM will automatically choose the sections which will
completely use your disk.

Number of
Sections                    Sections used

   1          The Section(s) used are 2
   2          The Section(s) used are 7,8
   3          The Section(s) used are 3,7,9
   4          The Section(s) used are 3,4,5,7
   5          The Section(s) used are 8,1,3,6,9
  Any         Choose your own sections
```

If the disk is *not new* (some sections were previously defined), only sections that have not previously been configured will appear in the listing. Place an x in the section(s) you wish to configure. As you choose sections the overlapping sections will be made unavailable. If you select the wrong section simply put the cursor on that section and press the Space bar.

8-10. SLIDE: Adding a New Disk to Series 800 Systems

Instructor Notes

Purpose

As we did earlier with the workstations, we want to show how a device and its sections can be selected for configuration into the system.

Key Points

- This screen is only to be used for adding a new disk to the system configuration. If a disk has already been defined on your system to contain swap space or file system sections, only the undefined sections will appear as "available" on this screen.
- The "*Add a Local File System*" screen is best used to build or rebuild file systems on selected disk sections whereas "*Add a Hard Disk*" is intended to define a whole disk and all of its sections at once (both file systems *and* swap). File systems can be added and built from either screen but you cannot define swap space from the "*Add a Local File System*" screen.

Teaching Tips

- Yes, this is almost the exact same material as the workstation section. If you have already presented that material, just note the different hardware addressing schemes and the fact that you have an extra step to select the disk sections.
- You should review the steps to physically add a hard disk here. This would be a good class review question.

  The steps for adding a disk are:

  1. Shut the system down and turn off the power.
  2. Choose the interface card to which the new disk will be connected.
  3. Choose a bus address for the new disk which does not match the bus address of any other device on the chosen interface card.
  4. Connect the new disk to the chosen interface card.
  5. Turn the new disk on.
  6. Turn the system on.
  7. Execute SAM to complete the process for adding the new disk.

Transition

We just selected the disk and sections to add now we need to define its usage to the system.

8-11. SLIDE: Defining a New Disk on the Series 800

Defining a New Disk on the Series 800

```
SAM                          Add a Hard Disk Drive

          Fill in or modify the desired fields and then press "Perform Task".
                       Card Slot Number      Bus Address      Interface Type

                              52                 5.8                scsi


 Initialize disk? (y or n) y
                                                        Create New      View/Modify
 Section     Size        Usage      Mount Directory     File System      Defaults

 7           75712       swap       ____________            n               n
 8           443528      hfs        /usr________            y               y
 10          129824      hfs        /users______            y               n
```

Help	Main Menu	Shell	Perform Task	hpterm		Disk Info	File Sys Info	Exit Task

50722 8-11 79

Student Notes

Once you have selected the disk to add and the sections you are configuring, you will see this screen.

Each line on the screen represents one of the sections you selected from the specified device.

The fields on the screen should be filled in as follows:

- *Initialize disk?* — Initializing a hard disk formats the physical media and tests it so it can be used by HP-UX. You should initialize a hard disk if the disk is new or has never had HP-UX on it before. You can only answer y to this question if you selected a combination of sections that comprise the whole disk. You cannot initialize just one section.
- *Usage* — This determines what the section will be used for. Your choices are:

  hfs Used for normal HP-UX hard disk file system section.

swap To define this as a dedicated swap section.

cdfs For CD-ROM file system.

ignore Used to define this section but leave it unused for now.

The choice you enter here will affect the rest of the fields you see.

- *Mount directory* — If you specified a usage of hfs or cdfs you must enter the mount point directory name for the file system. The directory must be empty or SAM will warn you that the files in the directory will be inaccessible while the file system is mounted.
- *Create a new file system?* — This only applies to hfs disk sections. Answering y to this question will cause SAM to build a new file system on the section. Any existing data on this section will be destroyed.
- *View/Modify Defaults?* — Answering n will accept the default values for these options. Answering y to this item will bring up a popup screen that allows you to specify several configuration options:

```
                     Default File System Options

      Fill in or modify the desired fields and then press "Done".

   Write protection (mark one) . . . . . .   read only   x read/write

   Set user ID execution allowed? (y or n)  y

   Mount . . . . . . . . . . . . . . . . .  x now       x on boot

   Long file names? . . . . . . . . . . .   y
```

The options you can change on this screen for hfs and cdfs sections are:

- *Write protection* — This allows you to select if the files and directories on the section may be changed or not. Marking a file system as read-only indicates that you cannot edit, create, remove or rename files on that file system. This works well for CD-ROM file systems because they are inherently read-only. Most file systems will be marked read-write. Read-write is the default choice.

- *Set user ID execution allowed* — A set-user-ID program is a special type of program that turns you into another user while the program is running. This allow you to specify whether users can run these special types of programs on the file system. If you are not sure about set-user-ID programs, mark the field with a y. Many of the system programs need this feature to operate correctly. The default choice is y.
- *Mount* — Selecting *now* will mount the file system immediately upon completion of the task. Selecting *on boot* will configure the system to automatically mount the file system at boot time.

hfs file systems can also be configured with "Long file names". The standard UNIX file name length is 14 characters. If you answer y to this question, this new file system will allow file names up to 255 characters long. If your applications can work on a long file name file system, you should answer y to this question. If you are not sure, answer n. If later you find that you would like to use long file names, you can convert the file system to support them. The default is to use long file names.

The only option you can change for swap sections is "Enable". Marking "now" will activate the swap section immediately upon completion of the task. "on boot" will configure the system to activate this swap section at boot time. Once a swap section is activated it cannot be deactivated without rebooting the system.

Fill in the fields as described then press `Perform Task` . You will see several status messages on the screen. When SAM completes the task it will ask if you want to add another disk. If you answer n you will be brought back to the "*Disk and Swap Configuration*" screen.

The defined disk is now a part of the system configuration. If you re-enter this screen to add another hard disk, the newly defined sections will not show up.

Changing a Disk's Sectioning

If you decide for some reason that you want to remove the disk or sections from the system or if you want to change the file system and swap layout, you will have to remove the disk from the system first then add it again with the new layout. Any file systems on the disk will, of course, be destroyed when you create the new section layout so you must make sure any file systems are backed up before starting this procedure.

8-11. SLIDE: Defining a New Disk on the Series 800

Instructor Notes

Purpose

To describe what needs to be entered to define the new disk.

Key Points

- SAM will create a file system and mount it if you choose.
- Once you define a disk in this fashion, you will not see it on this screen again unless it is explicitly removed. Using the `Remove a Hard Disk ...` option.
- The only valid section for a CD-ROM file system is section 2.

Teaching Tips

- Most of the fields are fairly obvious. You could describe most of them in very broad terms. The students can find the details in the notes.
- Go over the example sections on the slide. You may even want to draw a picture of the disk being configured to better illustrate the options.
- If you covered this same material for the workstation just note the differences in SAM's screen layout and the questions asked.

Transition

Since everyone has at least one file system in their configuration and most will have more than one, we should look at managing them next.

8-12. SLIDE: Managing File Systems

Managing File Systems

```
SAM                              File Systems

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


    Local (HFS, CD-ROM) File System Configuration:
          Add a Local File System ...
          Modify a Local File System ...
          Convert File System to Long File Name ...

    NFS (Network File System) Configuration -->

    Disk and Swap Configuration -->

    Status Information:
          View File System Information
          View Disk Space Information
```

Help	Main Menu	Shell	Select Item	hpterm				Previous Menu

50722 8-12 80

Student Notes

The slide shows the screen you will use when you have existing file systems and you would like to work with them.

Note that the rest of the material in this module is common to all HP-UX systems except where noted.

```
# sam (Return)

     File Systems ->
```

If you have a new disk, you will have to add the hard disk to your system configuration using `Add a Hard Disk ...` before it can be managed using this screen. As we have seen, part of adding a disk is to configure the disk's file system layout.

The options on this screen perform the following tasks:

- `Add a Local File System ...` — This option allows you to define a new file system in your configuration. This is useful if there is already a file system on a hard disk or CD-ROM and you would like to define how it will be used on your system and if it is automatically mounted (made available) at boot time. We will not use this option a great deal because our file systems will usually be built and defined when we add the hard disk using SAM's "*Add a Hard Disk*" screen.

- `Modify a Local File System ...` — Once a file system has been defined you can mount it or unmount it. Both of these tasks can be done using this option. This will also let you change whether the file system is automatically mounted at boot time.

- `Convert a File System to Long File Names ...` — If a file system was built using the standard UNIX 14 character file name restriction you can convert it to use long file names (up to 255 characters) using this option. Once a file system has been modified to support long file names it can not be changed back to the short file name mode.

- `NFS Configuration ->` — NFS is a LAN service that allows you to mount another machine's file system to your local file system. This option deals with configuring and managing Networked File Systems. We will not be discussing this option in this course because it relies on having the network properly set up and running.

- `Disk and Swap Configuration ->` — This is another way to get to the screen where we saw the `Add a Hard Disk ...` choice.

- `View File System Information`
- `View Disk Space Information` — These two options allow you to see information regarding the configured disks and file systems. We will cover these two options in detail.

8-12. SLIDE: Managing File Systems — Instructor Notes

Purpose

To show the ways SAM can help us manage file systems.

Key Points

- We will not be covering NFS. If you have a network set up in the classroom you may want to show interested students how to configure NFS at the end of the module. SAM can do everything necessary to use NFS except hook up the cables.
- The disk and swap configuration can also be accessed through the `Disk Drives ->` menu on the "*Peripheral Devices*" screen.
- There are only two times you would use `Add A Local File System ...` :
  - If you have a disk (like a CD) that contains a file system but that file system is not yet defined in your configuration. (The `Add a Hard Disk` option could also be used as described earlier.)
  - If you want to rebuild the file system on just one section of a disk (and thus destroying the data there).

Teaching Tips

- The details of each option are not extremely important here.
- Keep the discussion in terms of file systems instead of disks or disk sections throughout the rest of this module. This will help keep the presentation more general.

Transition

If we have more than one file system configured then we can make changes to its definition using the `Modify a Local File System ...` option.

8-13. SLIDE: Mounting and Unmounting File Systems

Mounting and Unmounting File Systems

```
SAM                    Modify a Local File System

        Modify the desired fields and then press "Perform Task".

                                                Mount              Protections
 Device File    Mount Directory      Usage    Now    on Boot     RO/RW      SUID

 /dev/dsk/6x0   /                    hfs      y        y          rw         y
 /dev/dsk/5x0   /users               hfs      n        y          rw         y
 /dev/dsk/2x0   /cdrom               cdfs     y        n          ro         y
```

Help	Main Menu	Shell	Perform Task	hpterm		Disk Info	File Sys Info	Exit Task

50722 8-13 81

Student Notes

SAM can mount and unmount file systems for us using the "*Modify a Local File System*" screen selected from SAM's "*File Systems*" screen.

This screen displays the definition of each file system in the configuration. Any fields that are underlined are items we can change in the file system's definition. The information on the screen is as follows:

- *Device File* – The name of the device file (disk or disk section) that contains the file system. Note that this cannot change.
- *Mount Directory* – The contents of the file system on the specified device will be attached to the file system tree under this directory. On the slide we can see that the `/users` directory and its contents are actually on file system that is separate from the root file system. You will seldom change the directory a file system is mounted to.

- *Usage* – This defines what type of file system exists on the device. The two things you might see here are **hfs** or **cdfs**.
- *Mount Now* – The file system is currently mounted if a "y" appears in this field. You can change the value of this field to either mount or unmount a file system.
- *Mount on Boot* – If a "y" appears in this field the file system will be automatically mounted at boot time. All file systems marked for mount on boot can also be mounted using the `mount -a` command. See "Related Commands" below.
- *RO/RW* – Read-Only or Read-Write protections. Most file systems will be marked `rw`.
- *SUID* – Specifies whether the file system was mounted to allow set-user-ID programs to execute.

Example

If you wanted to unmount the `/users` file system so you could do a maintenance check on it, you could enter "n" in the *Mount Now* field for `/users`. Later, when the file system check is completed, you enter this screen and type "y" under *Mount Now* and the disk will be remounted when you press `Perform Task`.

CD-ROMs

A CD-ROM file system (CDFS) is controlled the same way as any other hard disk file system except that it is always "read-only". On a Series 800, the file system always occupies section 2 (the whole CD).

CD-ROM file systems are *not* automatically mounted on boot up even if the *Mount on Boot* field is marked with a "y". This is because if the CD is not in the drive at boot time you would get an error. To mount your CD-ROM file system later you would enter this screen and mark the *Mount Now* field with a "y" and press `Perform Task`.

Related Commands

As we have mentioned, you can mount all defined file systems using the `mount -a` command. You can also see what file systems are mounted using `mount` with no options:

```
# mount
/ on /dev/dsk/6s0 read/write on Thu Jul  9 13:24:59 1992
/users on /dev/dsk/5s0 read/write on Thu Jul  9 13:28:41 1992
/cdrom on /dev/dsk/2s0 read only on Thu Jul  9 13:28:53 1992
```

To unmount a file system you can also use the `umount` command. You must specify the file system to be unmounted either by its device name or mount directory. This example will unmount the CD-ROM file system listed above:

```
# umount /cdrom
# mount
/ on /dev/dsk/6s0 read/write on Thu Jul  9 13:24:59 1992
/users on /dev/dsk/5s0 read/write on Thu Jul  9 13:28:41 1992
```

The command `umount -a` will unmount all file systems that are currently not in use.

Note

You cannot unmount a file system that is in use. The root file system is always in use and therefore can never be unmounted.

File systems containing active file system swap are also in use and may not be unmounted.

Reference

There are many more options to `mount` and `umount`. These are explained in Section 1 of the *HP-UX Reference Manual.*

Information about defined file systems is kept in a file called `/etc/checklist`. The format of this file is explained in Section 4 of the *HP-UX Reference Manual.*

8-13. SLIDE: Mounting and Unmounting File Systems

Instructor Notes

Purpose

To introduce the students to mounting and unmounting file systems and how to change a file system's definition.

Key Points

- CD-ROM file systems are not mounted at boot time even if they appear in the `/etc/checklist` file. This is actually because of an oversight in the `/etc/rc` script.

Teaching Tips

- Mention `mount` and `umount` but do not dwell on their options. You might also mention `bdf` here. Many people use it to see what is mounted instead of `mount` because the output is more readable.
- Stress again that only the superuser (`root`) can use `mount` and `umount`.

Transition

We can see a great deal about the current status of our file systems using SAM.

8-14. SLIDE: Viewing File System Information

Viewing File System Information

SAM View File System Information

Select a file system to view and then press "Return" or "Show Details".

Device File	Usage	Mount Directory	Mount/Enable Now	Mount/Enable Boot	Unused/Total (Kytes)
/dev/dsk/6s0	swap		y		54624
/dev/dsk/5s0	swap		y	y	54624
/dev/dsk/6s0	hfs	/	y	y	139398/339066
/dev/dsk/2s0	cdfs	/cdrom	y	n	669074
/dev/dsk/5s0	hfs	/users	y	n	130743/339066

— Press "Exit Window" when you are done. —

Help			Show Details	hpterm				Exit Window

50722 8-14 82

Student Notes

This screen lists the information all of our defined and mounted file systems and swap devices. (Swap devices are marked with a "Usage" of `swap`.)

The free space and total file system sizes are given in Kbytes (1024 byte units).

To see more information on any one of the file systems, move the cursor to that line and press `Show Details` .

```
                    Detailed File System/Disk Information

Device file:                /dev/dsk/6s0
Mount directory:            /
Usage:                      hfs

Mount/Enable now:           y                Protections ro/rw: rw
Mount/Enable on boot:       y                             suid:  y

Total size:                 339066 Kb        Card Slot:          0
Unused size:                139733 Kb        Bus Address:        6
File system swap size:      0      Kb

Long file names:            y


              — Press "Exit Window" when you are done. —
```

The size information displayed on this detail screen will not be displayed unless the file system is currently mounted.

Related Command

The bdf command will display total, used and free file system space of all mounted file systems.

```
# bdf
Filesystem           kbytes    used   avail capacity Mounted on
/dev/dsk/6s0         339066  164771  140388    54%   /
/dev/dsk/5s0         339066  174416  130743    57%   /users
/dev/dsk/2s0         669074  669074       0   100%   /cdrom
```

Because of the neat tabular format, many people prefer to use bdf instead of mount to see what file systems are mounted.

8-14. SLIDE: Viewing File System Information — Instructor Notes

Purpose

We need a way to see what is mounted or not and how to find how much space is available on our file systems. SAM and `bdf` both do that.

Key Points

- File system information is in the `/etc/checklist` file. This file and all currently mounted file systems (even if they are not in the `checklist` file) are displayed by SAM.
- The `kbytes` field in `bdf` is the total file system size. `used` plus `available` only add up to 90% of the total. Normal users cannot write to a file system that is 90% full. This reserved percentage is called the **minfree** amount and is held free due to performance reasons. When an hfs file system reaches 90% of its total capacity write performance drops by 50%!.
- The `capacity` field in `bdf` is the percentage of *usable* capacity (100% means the file system is 90% filled).

Teaching Tips

- You may want to ask students to run `bdf` and notice that used+available=90% of the total kbytes.

Transition

The other item left to cover in this module is swap space management.

8-15. SLIDE: Managing Swap

Managing Swap

```
SAM                      Disk and Swap Configuration

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

              Disk Configuration:
                   Add a Hard Disk Drive ...
                   Add a Floppy Disk Drive ...
                   Remove a Hard Disk Drive ...
                   Change a Hard Disk Drive Address ...

              Swap Configuration:
                   Add Device Swap ...
                   Modify Device Swap ...
                   Add File System Swap ...

              Status Information:
                   View File System Information
                   View Disk Space Information
```

Help	Main Menu	Shell	Select Item	hpterm				Previous Menu

50722 8-15 83

Student Notes

SAM has a small section for performing swap space management. There are two ways to get into the swap configuration section of SAM:

1. `# sam` (Return)

   `File Systems ->`

   `Disk and Swap Configuration ->`

2. `# sam` (Return)

   `Peripheral Devices ->`

`Disk Drives ->`

The swap configuration choices are listed under "*Swap Configuration*". The screens are almost identical for all HP-UX machines. As you might expect there are some minor differences due to the Series 800 disk sectioning.

The three swap configuration choices are:

- `Add Device Swap ...`
- `Modify Device Swap ...`
- `Add File System Swap ...`

We will discuss these options on the next two slides.

The `View Disk Space Information` option will allow us to see our swap space configuration.

8-15. SLIDE: Managing Swap

Instructor Notes

Purpose

Simple introduction to SAMs swap space management section.

Transition

Device swap is more common than file system swap so let's take a look at it first.

8-16. SLIDE: Managing Device Swap

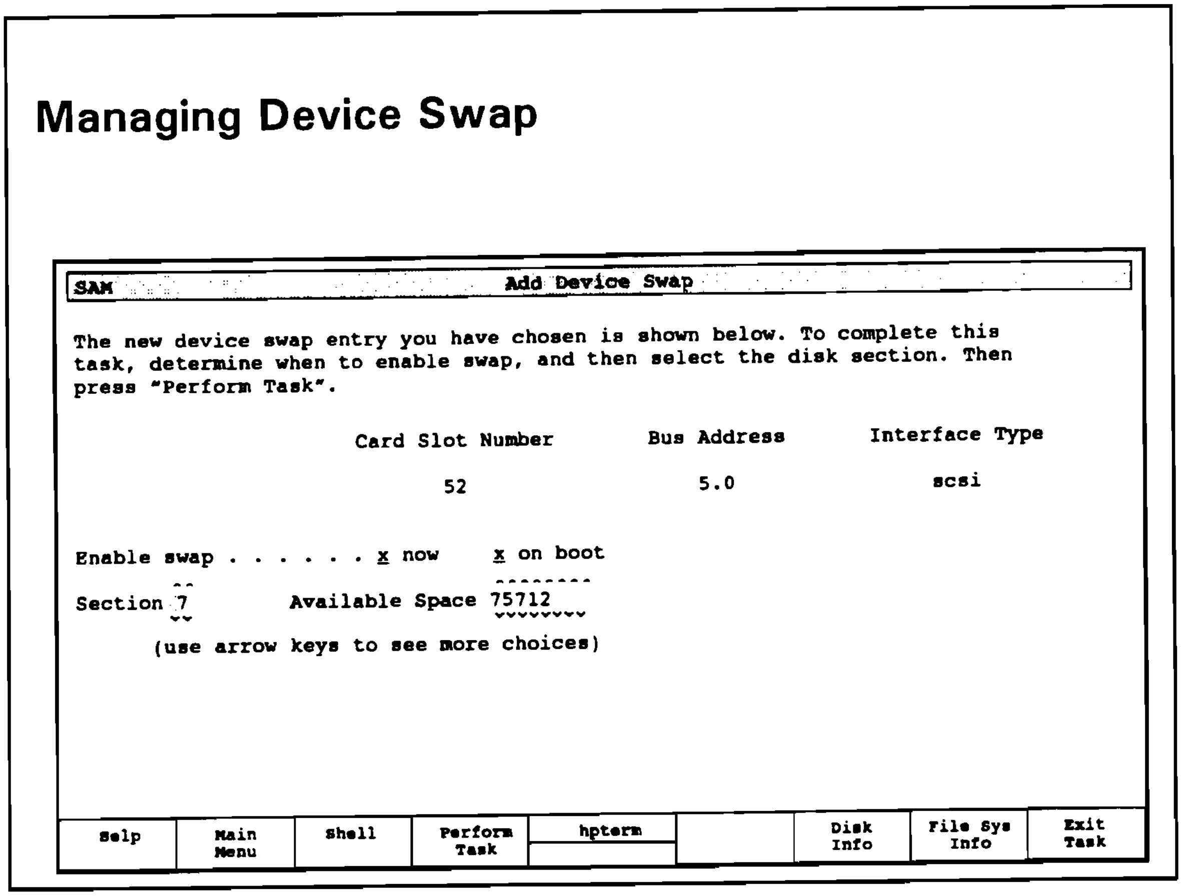

50722 8-16 84

Student Notes

You can define device swap space when you "*Add a Hard Disk*". This is the preferred way to define swap space. If at a later point in time you would like to add more device swap you can do so using the "*Add a Hard Disk*" screen again or the "*Add Device Swap*" screen. This slide is the "*Add Device Swap*" screen for a Series 800.

You can also use SAM to change your device swap configuration from the "*Modify Device Swap*" screen.

Adding Device Swap

Note

FOR WORKSTATION USERS: If you are working on a Series 300, 400 or 700 system and you add a disk with both file system and swap, the swap space will be configured at that time. You will not need to explicitly add it using SAM. You can use the "*Add Device Swap*" configuration screen to define swap space on a disk dedicated entirely to swapping (no file system).

Note

FOR SERIES 800 USERS: You can add device swap to any disk section that was not defined when you initially added the hard disk. Remember that you should only have one swap section per physical disk.

When you select `Add Device Swap ...` from the "*Disk and Swap Configuration*" screen SAM will prompt you for the device where the swap space will be configured. Of course, SAM will list the disks that have space available for swap and you simply select the correct one. If no unconfigured space is found, SAM will display an error window.

Once the disk has been selected you need to specify if the swap space you are defining will be enabled:

now which will activate the swap space immediately and/or

on boot which will automatically activate the space at boot time.

On the Series 800 you will also be able to choose the section where the swap space will reside. The available sections and their sizes will be listed on the screen for you. Simply select the correct section for your needs.

Once the fields have been filled in press `Perform Task`.

Modifying Device Swap

If at any point you want to change which swap devices will be activated at boot time, you can select `Modify Device Swap ...` from the "*Disk and Swap Configuration*" screen.

Note

You cannot deactivate swap space once it has been activated. If you no longer want to use a swap device you must mark it with an "n" in the "enable on boot" field then reboot your system.

You cannot change the parameters of the primary swap device. It is always activated at boot time even though its "enable on boot" field is often marked "n".

8-16. SLIDE: Managing Device Swap

Instructor Notes

Purpose

Most people will already have defined their swap devices when they added the hard disk. The purpose of this topic is to bring out the key points listed below.

Key Points

- There is no way to deactivate swap. Only the ability to keep it from being activated at boot.
- The primary swap is not marked for "enable on boot" because it typically does not get listed in `/etc/checklist`. It is activated by the kernel long before the root file system is even mounted.
- SAM is a little strange here. If you are on a Series 3/4/700 and you have a disk with a file system on it, you can not use `Add Device Swap` to define the space after the file system as swap. It expects that if you wanted to reserve swap space on that disk you would have done it when you did `Add a Hard Disk`.

Teaching Tips

- Note the differences between the workstations and the Series 800s. Generally stated, if you are administering a 3/4/700 system and your disks are all configured with file systems, you will never need to "*Add Device Swap*" because it was handled when you added the hard disk to the configuration.

Transition

File system swap is a way to add space to our swap configuration without dedicating a large portion of the disk.

8-17. SLIDE: Managing File System Swap

Managing File System Swap

```
SAM                        Add File Systaem Swap

File system swap is taken from the file system and should be used to augment
other swap rather than replace it. File system swap should not fill up the file
system. Move the cursor to the line of the file system you wish to change,
modify the desired fields and then press "Perform Task".

              Mount Directory    Minimum Swap    Maximum Swap   Total File System
                                   (Kbytes)        (Kbytes)       Size (Kbytes)

              /                    0_______        8_______         339066
              /users               0_______        20480___         339066

 Help | Main | Shell | Perform | hpterm |  | Disk | File Sys | Exit
      | Menu |       | Task    |        |  | Info | Info     | Task
```

50722 8-17 85

Student Notes

Recall that file system swap is intended only to supplement device swap, not replace it. File system swap is slower than device swap and your system will suffer a performance loss if the system uses the file system swap heavily.

When you select `Add File System Swap ...` from the "*Disk and Swap*" configuration screen, you will see a screen similar to the one on the slide. All currently mounted file systems will be listed.

To activate dynamic file system swap on any one of them simply fill in the appropriate fields:

- *Minimum Swap* specifies the amount of the file system the swap space will always occupy.
- *Maximum Swap* specifies the total amount of space that *may* be taken from the file system if the need exists.

Both of these values are given in Kbytes.

When you press `Perform Task` , a popup window appears and asks if you want this swap space activated now? If you are adding this file system swap for the first time you can answer y.

Another window will appear that asks if you want the swap space activated at boot time. If you would like to make this file system swap "permanent" answer y.

If you are modifying the minimum and maximum values at a later point in time, you cannot answer y to the activate now question. If you do and the swap space is already active on that file system, SAM will generate an error. It thinks you are trying to change the active swap space limits and that cannot be done.

To modify the minimum and maximum values correctly change the parameters on the "*Add File System Swap*" screen, then answer n to the activate now question and answer y to the activate on boot question. In this way SAM will not complain about changing an active swap system and it will save your changes so the correct swap parameters will be used the next time you boot.

An Example:

Here is a sample situation where file system swap works very well: The configured device swap on a system handles all of the swap requests under normal system usage. At the end of each month, however, several large applications are started that increase the need for swap space. Purchasing a new disk just for this end-of-the-month activity does not make sense. There are several large file systems with free space on them already configured into the system. Activating file system swap on these file systems does make sense. By specifying a small minimum amount, you will not be wasting file system space if the file system swap is not needed. Specifying a maximum amount that can accommodate the end-of-the-month activity will allow the system to use as much space as it needs dynamically.

8-17. SLIDE: Managing File System Swap — Instructor Notes

Purpose

To point out the intention of file system swap and how to configure it.

Key Points

- File system swap should not be a replacement for device swap.
- The points on attempting to modify active file system swap are important.

Transition

We can also look at the configured disks or disk sections.

8-18. SLIDE: Viewing Disk Space Information

Viewing Disk Space Information

View Disk Space Information

Select a disk to view and then press "Return" or "Show Details". If you are running on a cluster, only local disks are shown here.

Device File	Usage	Card Slot	Bus Address	Unused/Total (Kbytes)	File System Swap (Kbytes)
/dev/dsk/6s0	swap	0	6	54624	
/dev/dsk/5s0	swap	0	5	54624	
/dev/dsk/6s0	hfs	0	6	139235/339066	0
/dev/dsk/2s0	cdfs	0	2	669074	
/dev/dsk/5s0	hfs	0	5	130743/339066	0

—Press "Exit Window" when you are done.—

Help			Show Details	hpterm				Exit Window

Student Notes

This screen is similar to the "*View File System Information*" screen but this screen only contains information on the currently configured, locally attached disks. NFS mounted file systems are not included.

As you can see, one very useful piece of information on this screen is the hardware path to the device. This can be used to associate your defined disks to their device file names.

This screen has one other interesting piece of information, the "File System Swap" size. This screen is one of the easiest ways to see how much swap space is currently being taken out of the file system.

8-18. SLIDE: Viewing Disk Space Information — Instructor Notes

Purpose

This screen is especially useful for its hardware path and file system swap information.

Key Points

- This is the only place to easily see how much file system swap is currently being used on a file system.

8-19. WORKSESSION: Review Questions

1. What are three benefits of disk sectioning on a Series 800?

2. What are the required areas on a bootable disk on a Series 300/400? Series 700? Series 800?

3. What is the difference between the "*Add a Hard Disk*" screen and the "*Add a local file system*" screen in SAM?

4. What is the difference between the `View File System Information` option and the `View Disk Space Information` option on the "*File Systems*" screen?

8-19. WORKSESSION: Review Questions

Instructor Notes

1. What are three benefits of disk sectioning on a Series 800?

Answer:

1. smaller, more manageable file systems on large disks
2. better flexibility in file system layout
3. better control over file system usage

2. What are the required areas on a bootable disk on a Series 300/400? Series 700? Series 800?

Answer:

Series 300/400: 8Kb boot area, root file system, primary swap space.

Series 700: 8Kb LIF directory, root file system, primary swap space, 2Mb boot area.

Series 800: 2Mb boot area, root file system, primary swap space.

3. What is the difference between the "*Add a Hard Disk*" screen and the "*Add a local file system*" screen in SAM?

Answer:

"*Add a Hard Disk*" allows you to define a whole disk with either or both file system and swap space. "*Add a Local File System*" is only used to add an existing file system to the configuration. (On the Series 800 "*Add a Local File System*" will also let you rebuild a file system on the section(s)).

4. What is the difference between the `View File System Information` option and the `View Disk Space Information` option on the "*File Systems*" screen?

Answer:

`View File System Information` lists the mount point directory of the file system where the `View Disk Space Information` option lists the hardware address of the device and any file system swap information.

8-20. LAB: Exercises

Directions

Work with your lab partners to perform the following tasks.

1. If there is another disk available for your use, add that disk to your system configuration using SAM. Make sure you create a new file system on the device. Specify any mount point directory you would like. Make sure the disk is defined for automatic mounting at boot time. If you are on a Series 800, ask the instructor if there is a particular disk section you should use.

2. Use SAM to display the new file system information.

3. Unmount the file system you created and view the file system information again. What has changed?

4. Mount the new file system again and add some file system swap to it. Specify a minimum amount of 0 and whatever maximum amount you like.

5. Go into SAM's "*View Disk Space Information*" screen. Is your file system swap listed there?

6. Exit SAM then execute the `bdf` command and note what file systems are mounted.

8-20. LAB: Exercises

Instructor Notes

You should have more than one disk for each system to perform these exercises. If not you will have to coordinate the labs with available hardware so everyone gets a chance to try the exercises.

1. If there is another disk available for your use, add that disk to your system configuration using SAM. Make sure you create a new file system on the device. Specify any mount point directory you would like. Make sure the disk is defined for automatic mounting at boot time. If you are on a Series 800, ask the instructor if there is a particular disk section you should use.

Answer:

See the topic *Adding a Hard Disk* for your type of system for the complete details. The procedure will vary based on the system type used in class.

2. Use SAM to display the new file system information.

Answer:

```
# sam
    File Systems ->
  View File System Information
```

You should notice the device name and mount point of your new file system and all of its usage information. Also note that the "Mount Now" field is marked with a "y" meaning it is currently available for use. The "Mount on boot" entry should also be marked "y".

3. Unmount the file system you created and view the file system information again. What has changed?

Answer:

```
# sam
    File Systems ->
  Modify a Local File System ...
```

Mark the "Mount Now" field with a n and press `Perform Task` to unmount the file system.

If you `View File System Information` again, you will notice that the "Mount Now" field has changed to "n".

4. Mount the new file system again and add some file system swap to it. Specify a minimum amount of 0 and whatever maximum amount you like.

Answer:

```
# sam
   File Systems ->
 Modify a Local File System ...
```

Mark the "Mount Now" field with a y and press `Perform Task` to mount the file system.

Go into `Disk and Swap Configuration ->` then `Add File System Swap ...` to configure the swap space on the file system. In the "Maximum Swap" field enter the file system swap size in Kbytes. For example, to add 10Mb of file system swap you would enter 10240.

5. Go into SAM's "*View Disk Space Information*" screen. Is your file system swap listed there?

Answer:

```
# sam
   File Systems ->
 View Disk Space Information
```

The file system swap should be listed in the last column of the line for that file system.

6. Exit SAM then execute the `bdf` command and note what file systems are mounted.

Answer:

```
# bdf
Filesystem           kbytes    used   avail capacity Mounted on
/dev/dsk/6s0         339066  281934   23225    92%   /
/dev/dsk/5s0         386928  109538  238697    31%   /users
```

The exact output would vary with different systems. In this example there are two file systems: `/` and `/users`.

Module 9 — File System Maintenance

Objectives

Upon completion of this module, you will be able to do the following:

- Describe the file system superblock.
- Describe how HP-UX locates the control information and data for a file.
- Describe the relationship between directory entries and inodes.
- Explain how files are copied, moved, renamed, removed and linked.
- Identify the common causes of file system corruption.
- Explain how an HFS file system handles file modifications.
- Check for and repair file system corruption using `fsck`.
- Repair file system corruption.

About This Module

As system administrator, it is your responsibility to maintain the integrity of your file system. You can do that by checking it regularly for corruption, and repairing the file system if there are problems. This module teaches you how to do these tasks. This module also describes the basic HP-UX file system operation and structure so you can be better prepared to correct any possible file systems corruption.

Overview of Module 9

We want these students to know how to intelligently run `fsck` but they do not need to know the details of the HP-UX (hfs) file system structure and all of the details involved in tracing double indirect data pointers!

For this reason, this module starts out with an overview of the major structures of the HP-UX file system and how they work together to manage files. Once we complete this discussion, students should be able to effectively run `fsck` and understand what it is doing so they can answer its questions.

9-1. SLIDE: Understanding the HP-UX File system

Understanding the HP-UX File system

The HP-UX file system (HFS) is managed by many data structures:

- The Superblock
  - Contains all critical information about the file system (size, free space, etc.)
  - There is only one superblock per file system.
- The Directory Entry
  - Contains information like file and directory names and where these can be found on disk.
  - There is one directory entry for each unique file and directory name in the file system.
- The Inode
  - Contains all information about a file (owner, permissions, size, etc.)
  - There is one inode for each unique file and directory in the file system.

50722 9-1 87

Student Notes

The HP-UX file system, called HFS, is very complex. It is the kernel's job to keep track of all of the files and directories that are stored on a file system. The data structures listed on the slide are only three of the many different pieces of information that allow the kernel to accomplish its task.

On the next several pages we will discuss in some detail these three data structures and why they are important to your file systems. We will also discuss how HP-UX performs basic file manipulation routines like listing directories (`ls`), copying files (`cp`), moving and renaming files (`mv`), linking file (`ln`) and removing files (`rm`).

Once you understand how these data structures work and how the basic file system operations are accomplished, you will be well prepared to use the `fsck` utility to check for and repair the most common problems and types of corruption you will see with your HP-UX file systems.

9-1. SLIDE: Understanding the HP-UX File system

Instructor Notes

Purpose

This page introduces the module topics and the motivation for them.

Key Points

- The students will not become file system gurus but they will be able to intelligently look at `fsck` messages and know basically what it is looking for.

Transition

We will describe each of the major data structures over the next few slides.

9-2. SLIDE: The Superblock

The Superblock

The superblock contains:

- Critical information about the file system
- Static information which includes:
  - File system size
  - Block size
  - Fragment size
  - Disk characteristics
- Dynamic information which includes:
  - Total number of free data blocks
  - Total number of free inodes
  - File system clean flag

50722 9-2 88

Student Notes

HP-UX uses information in the superblock for various file system maintenance tasks.

The superblock is a contiguous 8Kb (8192 bytes) of disk space that describes the file system that resides on that disk or disk section. It is also used to keep track of the current state of the file system. If the superblock becomes corrupt, the filesystem becomes unusable.

Because the superblock is so important, HP-UX always keeps a copy of the superblock in main memory. The superblock on disk is updated from this copy periodically. HP-UX also keeps many copies of the superblock at various locations around the disk. If the primary superblock is lost, damaged, or becomes corrupted in some way, a copy can be retrieved from one of these locations. The redundant copies of the superblock contain only the *static* information that is listed on the slide. The dynamic information can be reconstructed with a command called `fsck` which we will examine shortly.

A list of the locations of the redundant superblocks can be found in a disk file created when the file system was made. The file name depends upon the type of system you have. On the Series 300/400/700, this file is /etc/sbtab. On the Series 800 it is called /etc/super_blocks. You should print a copy of this file and save it for future reference.

9-2. SLIDE: The Superblock Instructor Notes

Purpose

This slide sets up the importance of the superblock and the fact that it is duplicated in many disk locations. This will help us understand the need and use of `fsck -b`.

Key Points

- Discuss the contents of the superblock and reiterate its importance to the integrity of the file system. Most important is the static information. If the superblock becomes corrupt, the administrator can pick up the static information from a superblock copy and rebuild the dynamic information with `fsck`.
- Students should print `/etc/sbtab` or `/etc/super_blocks` so they can have a hard copy of the locations of the superblock copies. If it is the root disk that becomes corrupted you will not be able to retrieve the on-line copy.

Transition

The superblock contains all of the critical information about the whole file system. The directory entries and inodes contain all of the critical information about each individual file and directory.

9-3. SLIDE: Locating a File on Disk

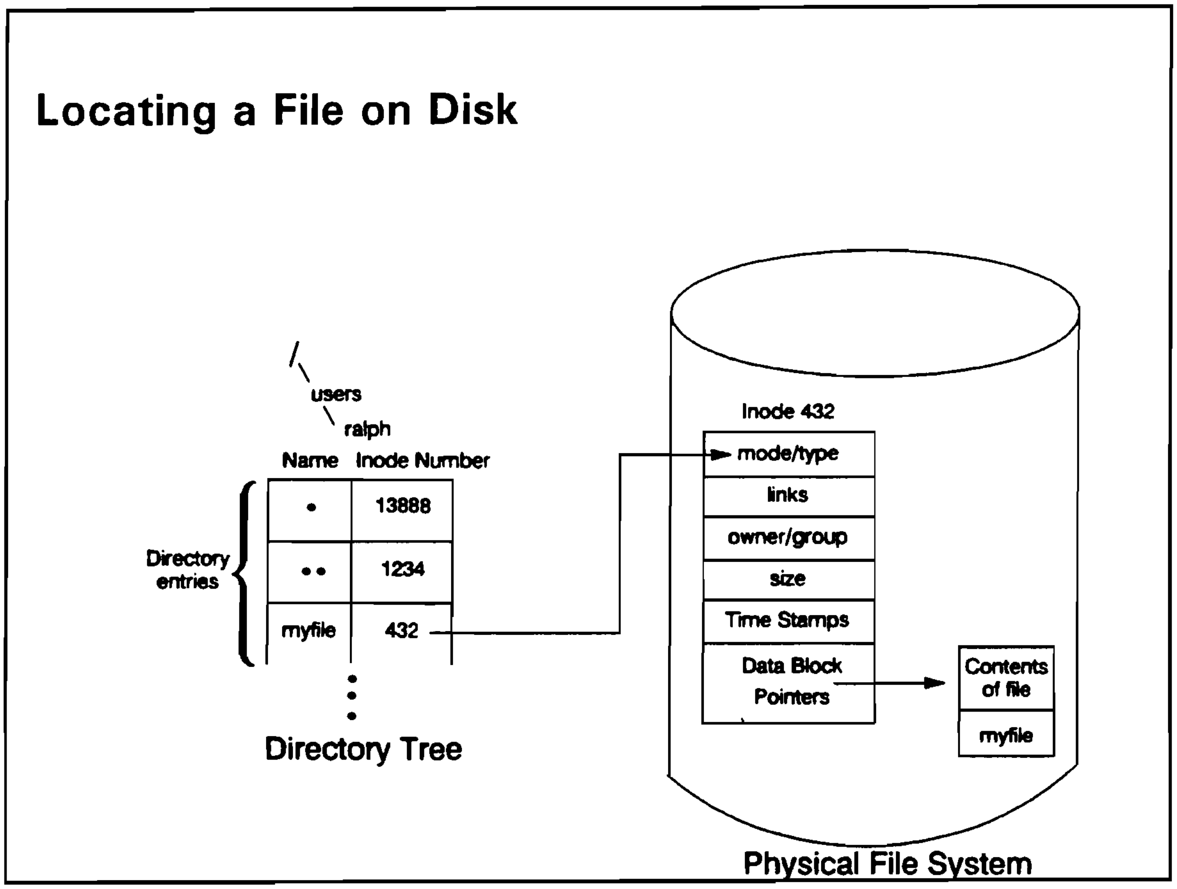

50722 9-3 89

Student Notes

We as users deal with the HP-UX file system in terms of file and directory names. The kernel has to know how to turn these logical names into physical locations on the disk. This task is accomplished using the three data structures listed on the slide.

Definition

Directory Entry. The directory entry of a file has two pieces of information: the file's name and its **inode number** (or **inum**). Every directory in the file system will contain multiple directory entries representing the files and directories stored there.

Inode. The inode (information node) is a physical location on the file system that contains all of the important information about the file. Every access to a file must first go through the inode. Inodes are referenced from the directory entries via the **inode number**.

Data Blocks. The pointers to the file's data blocks are also in the inode. The data blocks actually hold the file's contents.

On HP-UX, you can have a file system that supports either short filenames (up to 14 characters long) or long filenames (up to 255 characters). The maximum filename length is set when the file system is first created.

To see the inode number associated with each filename in the directory, you can use the `ls -i` command. The number to the left of the listing will be the actual inode number of this file. This can be useful if you are experimenting with commands like `cp`, `rm`, and `mv` and would like to see what these commands actually do to the file system.

The information stored in a file's **inode** is most easily seen using the `ls -l` or `ll` commands. Here are some details about the information in the inode:

File Type	Identifies the type of the file. This is either a regular file (first character is "-" in an `ll` listing), a directory (d) or some other special file type.
Mode	The mode is the permissions associated with the file. In an `ll` listing the mode is the set of nine characters identifying the "read" (r), "write" (w) and "execute" (x) permissions that apply to the file. (See chmod(1) in the *HP-UX Reference Manual.*)
Number of Links	The link count. The link count keeps track of how many filenames there are for a file. A file can have more than one filename by using the `ln` command.
Owner/Group	The ID of the user that created the file and the ID of the group the file belongs to. This information is used in conjunction with the mode to allow or deny access to the file.
File Size	The size of the file in bytes (characters).
Time stamps	These contain the time the file was last modified or read.

Example

In the example on the slide, there is a picture of a directory structure. If we tried to read `myfile` using the `cat` command, we would start a whole chain of events that use these data structures:

1. First, the corresponding inode number (432) would be retrieved from the directory entry and used as a pointer to inode 432 somewhere on the disk.

2. The ownership and permissions (mode) of `myfile` are retrieved from inode 432 to see if we are allowed to `cat` the file. To `cat` a file we must have "read" permissions on it.
3. If the requested operation is allowed, the data block pointer(s) of inode 432 would be traversed and finally lead to the actual data in `myfile` which would be displayed on the screen.

9-3. SLIDE: Locating a File on Disk — Instructor Notes

Purpose

This should help establish a high level perspective on HP-UX files and the role of inode numbers and inodes.

Key Points

- The three data structures work together to give us access to the files.
- We (the users) work on the file system in terms of directory and filenames. The system operates in terms of inodes and data blocks. The inode number in the directory entries is the link between the two.
- The file's name is stored only in the directory entry and NOT in the inode. We will see that this fact makes copying, linking and renaming files much easier.

Teaching Tips

- Make sure you cover the example that is on the slide using a simple command like `cat myfile`.
- A file's inode is only 128 bytes long. This is typically very small compared to the file's size. The slide is very disproportionate due to space limitations.

Transition

We now have a high-level understanding of how HP-UX manages files and directories. Let's take a look at some common file operations so we can see these data structures in action.

9-4. SLIDE: File Operations

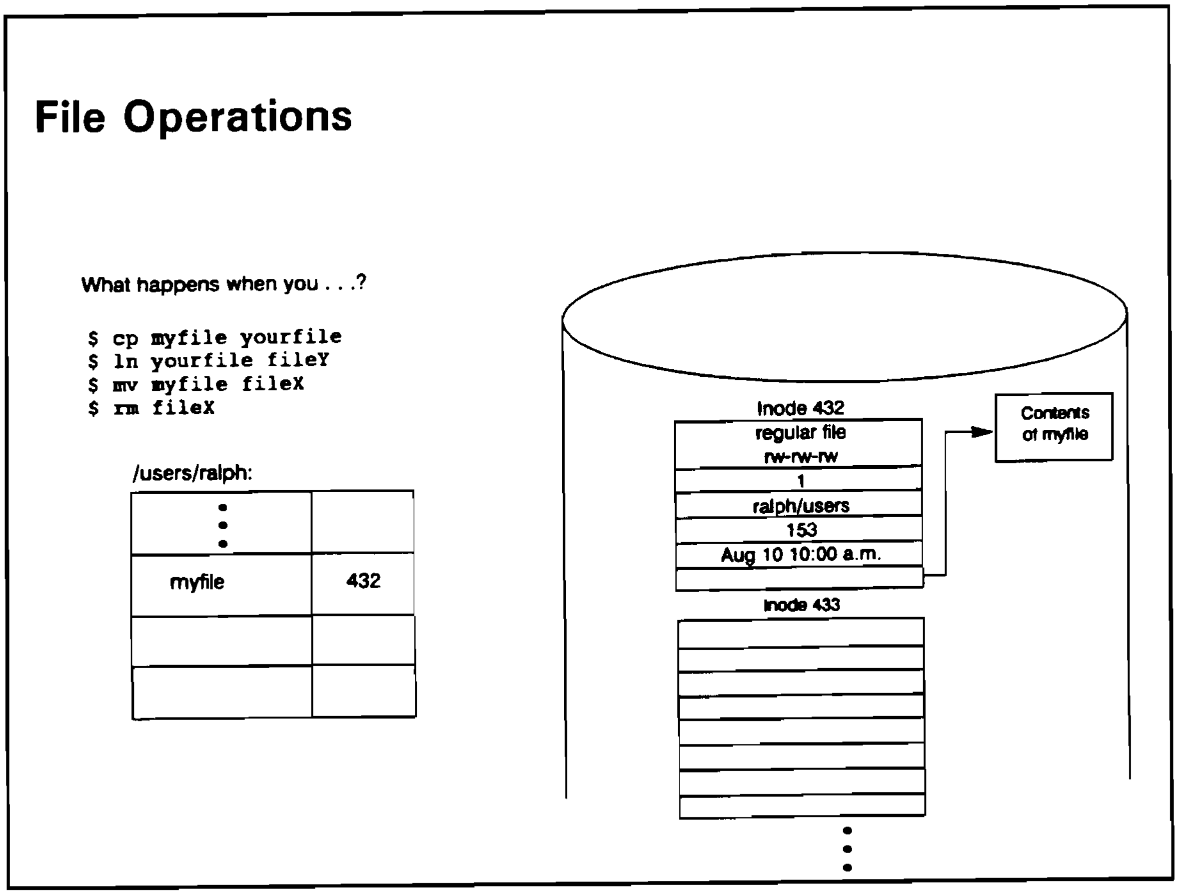

50722 9-4 | 90 |

Student Notes

We have all seen these file manipulation commands before. Have you ever thought about how they work?

What really happens when you ... ?

Copy a file

Using the cp command you make an exact duplicate of a file. Using the picture on the slide we can see what happens if you execute:

```
$ cp myfile yourfile
```

1. All of the information about `myfile` is read from its inode.

2. Since we are copying the data from `myfile`, another inode is allocated that contains the same size information (we can use inode 433 in the example on the slide).
3. `myfiles` data blocks are copied to new data blocks that the new inode will point to.
4. Another directory entry is also created that contains `yourfile` and the new inode number. The directory for yourfile could look like this on the slide: `yourfile - 433`

We now have two distinct names each referencing two different inodes. The files currently look the same but changing `myfile` will not affect `yourfile`.

Link two files

A link can be confusing if you are not sure how it works. It is really one of the simplest of the file system operations. When you link two filenames with the `ln` command like this:

```
$ ln yourfile fileY
```

You are simply creating a new directory entry for `fileY` with the same inode number as `yourfile`. The link count in the inode will be incremented to show that there are now two directory entries referencing this inode. In our example, inode 433 would now have a link count of 2. Doing an `ll` on `yourfile` or `fileY` will produce the exact same output showing that they actually refer to the same file:

```
$ ll -i yourfile fileY
433  -rw-rw-rw-   2 ralph    users      153  Aug 10 10:02 yourfile
433  -rw-rw-rw-   2 ralph    users      153  Aug 10 10:02 fileY
```

Now if `yourfile` is changed in any way, `fileY` is also changed because they reference the exact same data on the disk.

Rename or Move a file

The `mv` command is how we can rename a file.

```
$ mv myfile fileX
```

Since the filename only exists in the directory entry we simply have to change the name there. We do not have to modify the inode at all. Now the name in the directory entry for inode 432 on the slide would contain the name `fileX`.

If we are moving a file to a new directory like this:

```
$ mv fileX /tmp
```

we simply have to remove the existing directory entry and create another with the new name in the target directory (`/tmp` in this example).

Remove a file

To remove a file with `rm` we must mark the directory entry as being unused. This is done by replacing the inode number with a 0 (zero).

We must also decrement the link count in the inode since there is one less directory entry referencing this inode.

If the link count of the inode is now zero the inode will be returned to the pool of free inodes and the data blocks it pointed to will also be marked free.

If we were to

```
$ rm fileX
```

in our example on the slide, the directory entry for `fileX` would have an inode number of 0 and inode 432's link count would be decremented. Because its link count is now zero, inode 432 and its data blocks would be marked free for use by other files.

Link two directories

Note that it is very simple to link two files as described above, but because of the complexity in the file system traversal, we cannot link two directories in the same way. We must instead use a **symbolic link**:

```
$ ln -s /tmp/project /users/ralph/project
```

In this example, the directory `/tmp/project` already existed. The files and directories contained in it must also exist in `/users/ralph/project` (possibly because some test program expects it to be there).

Using the `ln -s` command we linked the two directories so doing any operations in the `/users/ralph/project` directory will be identical to doing those same operations in the directory `/tmp/project`.

Note that this is *NOT* like hard links where we simply create a new directory entry. A symbolic link actually creates a new file (file type `l` for link) that contains the pathname of the linked directory. In our example, when the kernel is trying to access `/users/ralph/project` it find that it is a symbolic link then goes to `/tmp/project` (where this link points and the files really live).

Symbolic links can also be used to link files or directories together that exist on different file systems. Hard links cannot cross file system boundaries.

9-4. SLIDE: File Operations — Instructor Notes

Purpose

To gain a high-level of understanding on how the most common file and directory operations are accomplished. Understanding these basic concepts and operations helps to understand the relationship between the directory entries (the file names) and the inodes (the files).

Key Points

- Symbolic links are not like hard links. You can actually remove the real directory that the link points to and get strange results. For this reason the use of symbolic links should be restricted to places where they are absolutely necessary.
- A file is not removed until its link count goes to 0.
- `mv` only changes directory entries, not inodes.

Teaching Tips

- Draw each example on the slide or on the whiteboard as you cover it. You must have pictures to go with this for it to really make sense.
- Encourage the students to experiment with the many variations of these commands. Most questions can be answered by trying it.
- If you have a very advanced class you may want to cover what happens if you copy or link to a file that already exists. In both cases the existing file's data is overwritten. In the case of a `cp` to an existing file, the new file retains the permissions of the file that was overwritten. Experiment!
- You may also want to mention that if you `cp` a file onto a file that is linked to another file, the link is maintained and the contents of the linked files will be the same as the `cp` source file. (The existing contents are overwritten). If you `mv` a file onto a linked file, the link is broken. Experiment with these too!

A Symbolic Link Example

Symbolic links can link a directory on one file system with one on another file system. This is a very common thing to do to balance the space usage between disks.

In the example, let's say that `/users` and `/tmp` are on two different file systems. If the filesystem containing `/users` was getting full we could decide to move some of the data to `/tmp`. We decide that ralph's `project` directory should go to `/tmp` because it is quite large. This is what we do to move the `project` directory to the `/tmp` file system:

1. Create the directory /tmp/project

```
# mkdir /tmp/project
```

2. Copy all of /users/ralph/project to /tmp/project

```
# cd /users/ralph/project
# find . | cpio -pdumv /tmp/project
```

3. Remove the entire /users/ralph/project directory

```
# cd ..
# rm -r project
```

4. Create the symbolic link as shown in the student notes.

```
# ln -s /users/ralph/project /tmp/project
```

Now the space is free on the /users file system and as far as the users are concerned the data still exists in /users/ralph/project (but it is really in /tmp/project).

Preparation Notes

Make sure you read *and understand* the manual pages for these commands. If you do not fully understand how these concepts work practice performing these file operations and use the ll -i and cat commands extensively between commands to help you see the relationships between the directory entries and the inode numbers.

Transition

Proper file system maintenance will help keep all of the filesystem data structures consistent so we will not loose any of our data stored in files.

9-5. SLIDE: File System Maintenance

File System Maintenance

- The system administrator is responsible for maintaining the integrity of the file system
- Corruption must be detected and repaired
- Maintenance tasks that should be performed regularly
  - Use `fsck` regularly
  - Ensure `sync` executes regularly
  - Monitor disk usage
  - Employ regular backup procedures
  - Always use proper shutdown procedures

50722 9-5 91

Student Notes

One of the principal responsibilities of a system administrator is to prevent the loss of the users' data. Since the data storage structure utilized by HP-UX is the file system, it's imperative that the file system be checked regularly for possible problems. We have seen that the file system is a fairly complex set of data structures. If these data structures are not absolutely consistent with each other, there is a very good chance that you could start loosing critical data.

There are many things the administrator can do to help maintain the integrity of the file systems, including those listed on the slide.

For the remainder of this module we will discuss the most important of these tasks, the `fsck` command. We will also discuss why proper system shutdown is so important to the file system and how the `sync` command works to keep the file systems up to date.

We have already covered how to properly shutdown an HP-UX system and we will cover system backups shortly.

9-5. SLIDE: File System Maintenance — Instructor Notes

Purpose

This gives a broad overview of the file system maintenance tasks.

Teaching Tips

- Do not get into details here. All of these tasks will be covered if they were not covered already.
- You may want to note that even normal operations on a file system can cause minor problems. That is why `fsck` should be run periodically.

Transition

The file system buffers give us a big performance boost but they also allow the biggest opportunity for file system corruption.

9-6. SLIDE: The Buffer Cache

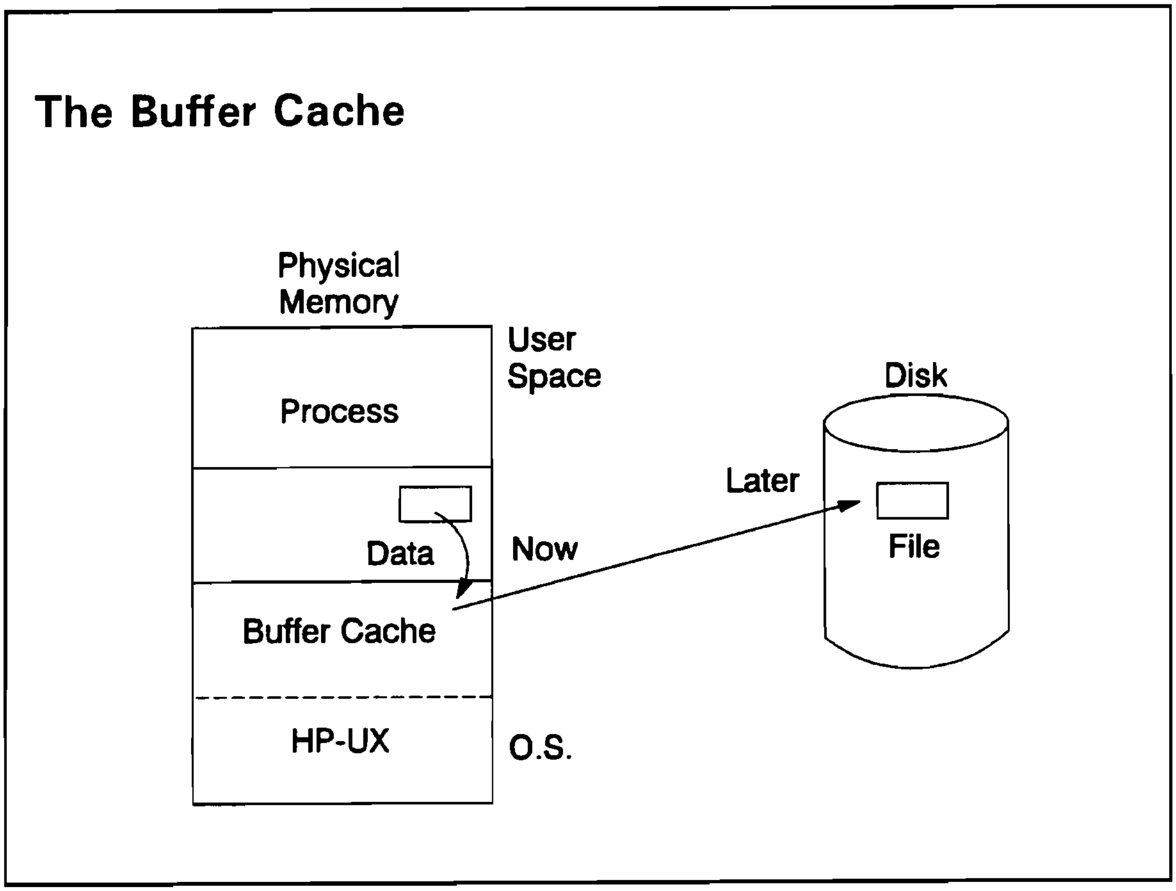

50722 9-6 92 ©1992 Hewlett-Packard Co.

Student Notes

When data needs to be written to a disk, the actual write does not occur immediately. The data is initially copied to an in-memory file system buffer called the **buffer cache**. This is a much faster operation than if the system had to perform an actual write to the disk. Your program continues satisfied that it has completed the write. The data from the buffer cache, along with any directory or inode information, is actually written to the disk sometime later. This is designed to improve overall file system throughput.

Note

If the system is halted without writing the buffers to disk, the file system could be corrupted.

In fact, improper shutdown is one of the most common causes of file system corruption.

Corruption due to an improper shutdown may be as trivial as having some incorrect counts in the superblock or as serious as corrupting a whole directory so you can no longer access the files there. In either case the problem is fairly easily fixed using the `fsck` command.

9-6. SLIDE: The Buffer Cache — Instructor Notes

Purpose

To explain and illustrate the purpose of the buffer cache, how it works and why it is very important to properly shutdown the system.

Preparation Notes

Chapter 8 "The HFS File System" in *How HP-UX Works: Concepts for the System Administrator* has good information about the file system and a little on the buffer cache.

9-7. SLIDE: The sync Command

The sync Command

- Writes buffer contents to disk
- Keeps the file system current
- Is usually invoked on a regular basis by the syncer program
  - `syncer` is executed in the `/etc/rc` file
  - The syntax of the syncer program is:

```
syncer [seconds]
```

50722 9-7 93

Student Notes

As we have seen, data is written to buffer cache before it is written to disk. The physical write from the buffer to disk is delayed until:

- The system needs the buffer for another operation.
- The file system is unmounted.
- The `sync` command is executed.

When the `sync` command is executed, it causes the system to flush its buffers and write all data to disk.

Whenever the system is in a multi-user mode there will be a daemon called `syncer` running. It was started by the `/etc/rc` script when the system first entered multi-user mode. The `syncer` keeps the file system current.

9-7. SLIDE: The sync Command

Instructor Notes

Purpose

To explain the purpose of the sync command and the syncer program.

Transition

sync minimizes data loss due to improper shutdowns. There are other types of corruption, however, where sync will not help much.

9-8. SLIDE: Causes of File System Corruption

Causes of File System Corruption

- Improper shutdown procedures
  - Not using `reboot` or `shutdown` to bring down the system
  - Taking a mounted file system off-line
- Improper startup procedures
  - Not checking a file system for inconsistencies
  - Not repairing inconsistencies found
- Hardware failure

50722 9-8 94

Student Notes

A file system can become corrupt in a number of ways. The most common way is improper shutdown procedures due to shutting off a running system or power failures in the facility. Even the day-to-day operations on a file system can cause minor inconsistencies on a file system. That is why it is important to run `fsck` periodically.

When the system boots up it checks all of its defined file systems to make sure they were properly shutdown. If they were not, they are checked for inconsistencies. You should never attempt to bypass this check. Running a file system that may have inconsistencies is very dangerous. Even small problems could eventually lead to data loss if left uncorrected.

Note

Just because a file system is found to be clean on boot up does not mean it is corruption free. It simply means the file system was properly shutdown or unmounted.

The hardware of a computer system is usually very reliable. However, any piece of hardware could fail at any time. Hardware failures vary and can be as subtle as a bad block on a disk pack or as obvious as a non-functional disk controller. By following recommended hardware preventive maintenance procedures you can avoid the most serious problems and be prepared for any that might occur.

Regardless of the reason for the corruption you can minimize your risk of data loss by running `fsck` periodically and keeping regular backups.

9-8. SLIDE: Causes of File System Corruption — Instructor Notes

Purpose

To explain causes of file system corruption. By knowing what causes file system problems we can avoid it.

Key Points

- Improper shutdown is the primary cause of file system corruption.
- `fsck` should be run periodically to find inconsistencies even on disks that were "clean" on boot up.

Transition

Finally, we can talk about the tool to check and correct file system problems.

9-9. SLIDE: The `fsck` Command

The fsck Command

- Checks file system consistency and makes repairs
- Is multi-pass
- Should be run on an unmounted (or quiet) file system
- Can be run interactively or non-interactively
- Syntax:

```
fsck -p|-P [file_system]
fsck [-b block#] [-y | -n] [-q] [file_system]
```

50722 9-9 95

Student Notes

`fsck` (file system check) is the principal file system maintenance tool available with HP-UX. It verifies the structural integrity by checking data which is intrinsically redundant in a file system. The redundant data is either read from the file system or computed from known values.

`fsck` is a multi-pass program, meaning that it examines the file system a number of times, each iteration examining a different feature of the file system. Each pass `fsck` makes through the file system is known as a **phase**.

During each phase, any inconsistencies noted in the file system are reported. If you are running `fsck` interactively it will also ask if corrective action should be taken. Answering y (for yes), will always tell `fsck` to correct the problem. A n (no) response, tells `fsck` to ignore the inconsistency found and continues its checking. There are very few occasions when a "no" response should be given. If you are running `fsck` non-interactively, it will fix all of the most common file system problems without asking questions. If an inconsistency can only be fixed by removing a file's data, the non-interactive `fsck` will stop and tell you to run `fsck` interactively. We will discuss both forms of the command.

`fsck` should be run on quiet file systems. Ideally, the file system should be unmounted. Since this is not possible for the root file system, you should bring your system to a quiet single-user run-level with the `shutdown` command before running `fsck`.

The *file_system* argument is optional. If you invoke `fsck` without giving it a *file_system*, `fsck` will check the file systems defined in the `/etc/checklist` file. You can check just one file system by supplying the device file name for the file system on the command line.

`fsck` does not actually correct any of the problems it finds until the very end of the check. If you want to abort the check for some reason you can do so by typing Ctrl C any time before `fsck` reaches its final phase. The file system will be left unchanged and you *should run* `fsck` *on that file system later.*

Action

We have said that `fsck` should be run on a regular basis to keep small problems from turning into data loss.

For preventive maintenance, it is a good idea to run `fsck` on all file systems about once a week.

Reference

The *HP-UX Reference Manual* contains the details of the `fsck` command.

The *Error Message Catalog* contains a whole section on the all of the errors and questions `fsck` may display. It describes their meaning and possible implications of your response.

9-9. SLIDE: The `fsck` Command — Instructor Notes

Purpose

To introduce the purpose and use of the `fsck` command.

Key Points

- A common problem is to have a data block listed as being free when in reality it is already allocated to a file. This problem is easily found and corrected by `fsck` if you run it regularly. It simply marks the data block as used. If this problem goes unchecked, however, another inode might actually be assigned the same data block (because it was marked free). Now we have lost data because the situation cannot be remedied without removing the older file.
- When you get an error you don't understand from `fsck`, look it up in the *Error Message Catalog*. If it is still not obvious what `fsck` is talking about you can just answer y to fix the problem and potentially remove data. The safer choice is to call an HP representative or the Response Center for help.
- The *file_system* name is typically the block device file name of the file system. You can supply the character device file name only if the file system is unmounted. Since the root filesystem is always mounted you should *never* check it with the character device file.

Transition

We will talk about both the non-interactive and interactive forms of the `fsck` command on the next two slides.

9-10. SLIDE: fsck in Preening Mode

fsck in Preening Mode

```
fsck -p [ file_system ]
fsck -P [ file_system ]
```

Non-interactive mode

Fixes inconsistencies but does not remove data

Will correct:

- Unreferenced inodes — files with no directory entry
- Link counts that are too large
- Missing blocks in the free list
- Blocks in the free list also in files
- Incorrect counts in the super-block

50722 9-10 96

Student Notes

The non-interactive invocation of fsck is called preening mode. Both the -p and -P options run fsck in this mode.

This option corrects many problems, but *never removes data*. For each problem it corrects, it prints a message identifying the file system and the corrective action taken. If it cannot solve a problem, fsck terminates and tells you to run it in interactive mode. If fsck terminates in this way, you must run it interactively so you can work with fsck to correct the problem.

The -P (upper-case) option operates in the same manner as the -p (lower-case) option except that those file systems which were cleanly unmounted will not be checked. (fsck checks the clean byte of the superblock.) fsck -p will check the file system non-interactively regardless if the file system was properly shutdown or not. fsck -P is the command which is run at bootup time. It is invoked in the /etc/bcheckrc script.

9-10. SLIDE: `fsck` in Preening Mode — Instructor Notes

Purpose

To explain how to run `fsck` in preening mode.

Key Points

- The normal boot-up mode for `fsck` is `fsck -P`.
- Preening mode will fix any inconsistency that will not result in data loss, or it will abort (describing the problem) and ask the administrator to run `fsck` interactively if it finds an inconsistency that could result in data loss.

Teaching Tips

- If the group did well with the beginning of this chapter of the file system structures, you may want to ask what each of the items listed on the slide really mean and why they are a problem. For example, an unreferenced inode is an allocated inode that has no directory entry pointing to it. A link count that is too large means the file's data will never be removed from the system because all of the directory entries will be removed and the link count will not go to 0. If all of the directory entries are indeed removed you will have an unreferenced inode. We have already discussed the problem with having blocks in the free list also allocated to files. Having incorrect counts in the superblock can make you and the system think there is more data space free on the disk than is actually free. This can be a problem if the file system is getting full.

Transition

The interactive mode of `fsck` has many other options.

9-11. SLIDE: Other Options to fsck

Other Options to fsck

```
fsck [ -b block# ] [ -y | -n ] [ -q ] [ file_system ]
```

- `-b` Uses the block specified immediately after the flag as the super block for the file system
- `-y` Assumes a "yes" response to all questions asked by `fsck`
- `-n` Assumes a "no" response to all questions asked by `fsck`
- `-q` Quiet mode

50722 9-11 97

Student Notes

If you invoke `fsck` with no options or the options listed on the slide, it runs interactively. In interactive mode, `fsck` poses a question when any inconsistency is found and waits for your response. As discussed earlier, a `y` response will fix the problem and an `n` response will ignore the problem. In either case `fsck` continues searching for the next problem.

The annoying thing about `fsck` when used with no options is that it asks you questions about inconsistencies that are easily corrected (like those listed on the previous slide).

The options available to `fsck` shown on the slide can modify the behavior of `fsck` so it works more automatically. These options are described in more detail here:

`-q` Quiet Mode. This is the best combination of interactive and non-interactive modes. It can get quite annoying answering all of the questions `fsck` will ask even for problems that you are sure should be fixed. The `-q` option will fix all of the problems described in the preening mode section but if it encounters a serious error that may require removing data, it will stop and ask you the appropriate question. This is the recommended way to check your file systems interactively.

`-n` This option causes `fsck` to answer "no" to all questions posed by the command. Since this option never results in the loss of data, it may appear safe. However, since "no" is supplied as the answer to all questions, `fsck` takes no corrective actions and inconsistencies are not resolved. It is recommended that this option be used when you just want to assess the state of the file system; however, you should invoke the command again with a different option to resolve inconsistencies if any are found.

Example:

```
# fsck -n /dev/dsk/c0d0s3  |  tee /tmp/fsck.log
```

In this example, the diagnostic output is directed to the file `/tmp/fsck.log` as well as to the screen. This output can be analyzed later to determine corrective action. If you see errors in the log file that you do not understand, look them up in the *Error Message Catalog* to determine how to handle them. Of course, `/dev/dsk/c0d0s3` is only an example file system device name. You would enter the device file name for your file system instead.

`-y` This option causes `fsck` to answer "yes" to all questions posed by the command. It is possible that data will be removed as a result of a "yes" answer. Consequently, if this option is to used, the file system should be examined with the `-n` option first so the possible consequences can be assessed.

`-b` *block#* If you get a superblock or "unrecognized file system" error from the `mount` or `fsck` commands, the primary superblock is probably corrupt. The `-b` option tells `fsck` to copy the superblock at *block#* into the primary superblock then do the file system check. You can find the superblock copy numbers in `/etc/sbtab` on a Series 300/400/700 or `/etc/super_blocks` on a Series 800.

Note

If you are running `fsck` on a mounted file system (like the root file system) `fsck` may ask you to reboot the system using `reboot -n`. This will keep bad data in the cache from being flushed onto the clean file system.

If `fsck` gives you a message like this you must reboot using `reboot -n` immediately. Continuing to run the system could cause the system to become corrupted again.

9-11. SLIDE: Other Options to fsck

Instructor Notes

Purpose

To describe the other important options of `fsck`.

Key Points

- Stress that most of the time, if `fsck` is run periodically, it will do its work and find no problems. If problems are found they will be usually simple to fix.
- `-q` is one of the most useful interactive options.
- If your file system has a corrupted superblock, `fsck` will tell you that there is a problem and you will have to correct it with `fsck -b`. You may also see superblock errors when using `mount`.
- The `-n` option fixes nothing. It is just used to see what `fsck` finds.
- If you are using a standard 8Kb data block file system the first superblock copy is located at block number 16.
- If the root file system's superblock is corrupt, you could be in trouble. You will most likely have to fix the disk by booting from another disk or a recovery system and running `fsck` on your real root file system. This is a very complicated problem and is avoided here.

Preparation Notes

You will have to run `fsck` in a variety of situations in order to really understand how it works and what type of questions it asks. To practice, try using the `clri` command to invalidate an inode on a file system then run `fsck` on that filesystem. You should get at least 3 errors from the different phases. You could also try to `fsck` the root file system after using `clri` on an inode. This should generate the message telling you to use `reboot -n`.

A Common "Serious" Problem

One of the most common "serious" `fsck` problems is when a data block is referenced by more than one inode. This is the result of the block being allocated to one inode but still marked free. The second inode then also claimed that space. The only way to remedy this problem is to remove the original (older) file's other data blocks (the newer file has overwritten the data of the older file). `fsck` calls this a duplicate block error and asks if you want to clear the older file. The problem is that you don't know what the older file's name is. All you have is an inumber and some size and ownership information that `fsck` reported.

At this point, if you want to find the file's name you must stop `fsck` using Ctrl C. Using the inumber of the older file that `fsck` gave, you can find the file name using the `find` command:

```
# find dir -inum inumber
```

Note that the file system in question must be mounted for this to work. *dir* is the mount point directory for the file system in question (this will be / for the root file system).

Once you know the name you can safely clear the file using `fsck` then restore the file from a backup at a later time.

Transition

`fsck` needs a place to put the orphaned files it finds. That place is the `lost+found` directory.

9-12. SLIDE: The lost+found Directory

The lost + found Directory

- At the top of every file system
- Created when the file system is created
- Must have some empty file slots
- fsck copies problem or orphaned files to this directory
- Check this directory after each invocation of fsck and try to determine the origin of any files there
- Can be created with the mklost+found command

50722 9-12 98

Student Notes

Every file system should have a lost+found directory at the root of its file system (for example, the one for the root file system is called /lost+found). The lost+found directory is created Automatically when you create the file system. However, you should verify that the directory exists before using fsck to check the file system. If lost+found does not exist, you can rebuild it with the /etc/mklost+found command.

The lost+found directory has empty directory entries so that fsck can write to it without allocating more disk space and possibly changing the disk's block counts.

fsck places any problem files or directories in the lost+found directory. These files and directories are typically the "unreferenced inodes" that we discussed earlier. fsck links them into this directory so that we can reference them and possibly put them back where they belong.

After `fsck` completes, you should examine the contents of the directory. The files that are placed there should be moved back to their original directories. It may be difficult (and sometimes impossible) to determine where the files actually belong. Doing an `ll` on the `lost+found` directory will show you the file type, permissions, ownership, size and modification time but the name will be a number assigned to the file by `fsck`. If you know who owns the file, you can ask them to put it back where it belongs. If the file is owned by a common system user like `root` or `bin` you will have to do more work to determine what the file contains. There are several tools you can use to do this:

`more` *filename*	If the file is a text file, `more` will display its contents to your screen.
`file` *filename*	Display the file's type.
`strings` *filename*	`strings` displays any text strings the file may contain (like error messages or revision control strings). The name of the program will often appear in one of the strings.

Note

The filenames in `lost+found` are numbers with a leading `#`. The `#` character denotes a comment in the shells. Thus, if you tried to `more #01234` you will get a usage error.

You must use a \ before the `#` to make it work correctly on the command line:

```
# more \#01234
```

9-12. SLIDE: The lost+found Directory — Instructor Notes

Purpose

To introduce the lost+found directory and some tools to find where the orphaned files belong.

Key Points

- The leading # on filenames in lost+found are often the source of some confusion when trying to work with these files.

Teaching Tip

- Have the students try some of the tools mentioned on known files like /bin/cat and /etc/passwd. This will let them try these commands and see the output.

Transition

It's time to try some of the things we have learned about file system maintenance.

9-13. WORKSESSION: Review Questions

1. Briefly describe the buffer cache.

2. What one disadvantage of the buffer cache does the System Administrator have to worry about?

3. When does `fsck` actually fix the disk?

4. How can a system administrator prevent having a corrupt file system?

5. What is the `lost+found` directory used for?

9-13. WORKSESSION: Review Questions

Instructor Notes

1. Briefly describe the buffer cache.

Answer:

The buffer cache is an area of physical memory that is used as a buffer between user memory space and disk space. This cache enhances overall file system throughput.

2. What one disadvantage of the buffer cache does the System Administrator have to worry about?

Answer:

The fact that the buffers need to be flushed to disk before shutting down the system. If not, `fsck` needs to be run to check the file system that may have been corrupted.

3. When does `fsck` actually fix the disk?

Answer:

`fsck` checks the file system and notes all necessary changes. The problems are not actually fixed until the very last few seconds that `fsck` runs. For this reason you can Ctrl C out of `fsck` at most any time and the file system will not be changed.

4. How can a system administrator prevent having a corrupt file system?

Answer:

By using the proper shutdown and startup procedures and running `fsck` periodically (about once a week).

5. What is the `lost+found` directory used for?

Answer:

It is used by `fsck` to store orphaned (unreferenced) files. The only problem with the file is that it had no directory entry. `fsck` makes a directory entry for the file in `lost+found`.

9-14. LAB: Exercises

Directions

Work with your lab partner(s) to perform the following exercises.

1. Use the `ls -ai` command to find the inode numbers associated with the entries in the root directory.

2. Create a file called `test.file` with some lines in it. Use the `ls -i` command to find its inode number.

3. Now, rename the file to your name (use the `mv` command). Do another `ls -i`. Has its inode number changed?

4. Make a copy of your new file into `fileX` (using `cp`). Does the files have the same inode number now?

5. Hard link another file to `fileX` and call it `fileY`. Use `ll -i` to see the inode number of `fileY`. What happens to `fileY` if you remove `fileX`?

6. On which file systems would the command `fsck -p` perform a file system check if you did not specify a file system on the command line?

7. If applicable, run `fsck` on the secondary file system you created and mounted in the module on creating file systems. Use the preening mode of fsck. Note any problems you may encounter and how to fix them.

8. Properly shutdown your system and run `fsck` on your root file system disk. Again, note any problems you may encounter.

9-14. LAB: Exercises

Instructor Notes

Warning **Be sure that students do NOT have the filesystem swap enabeled from their previous lab before attempting this lab, otherwise they will not be able to unmount their filesystems!**

Directions

One way to corrupt a secondary file system is to use the `clri` command. Assuming /disc is the mount point for the secondary file system we can corrupt the file system like this:

```
# touch /disc/junk
# ls -i /disc/junk
# umount /disc
# clri /dev/dsk/device_file_name inumber_of_junk
```

Some notes on `clri`: For proper results, the file system should not be mounted. After `clri` is executed, the inode has no valid information in it and any blocks in the affected file show up as "missing" in an `fsck` of the file system.

If the file system *is* mounted when you do `clri`, the inode information in memory may simply rewrite the inode and you did nothing. If the inode was not in cache, you run the risk of crashing your system if you try to access the cleared inode!

1. Use the `ls -ai` command to find the inode numbers associated with the entries in the root directory.

Answer:

```
# ls -ai /
```

2. Create a file called `test.file` with some lines in it. Use the `ls -i` command to find its inode number.

Answer:

```
# vi test.file
# ls -i test.file
```

3. Now, rename the file to your name (use the `mv` command). Do another `ls -i`. Has its inode number changed?

Answer:

```
# mv test.file your_name
# ls -i your_name
```

The inode number of *your_name* will be the same as `test.file`.

4. Make a copy of your new file into `fileX` (using `cp`). Does the files have the same inode number now?

Answer:

```
# cp your_name fileX
# ls -i your_name fileX
```

The two files will not have the same inode numbers because they are two distinct copies.

5. Hard link another file to `fileX` and call it `fileY`. Use `ll -i` to see the inode number of `fileY`. What happens to `fileY` if you remove `fileX`?

Answer:

```
# ln fileX fileY
# ls -i file?
```

The inode number is the same for all files that are linked together. Notice also that the link count is 2.

Removing `fileX` will have no effect on `fileY` except that the link count will decrease to one.

6. On which file systems would the command `fsck -p` perform a file system check if you did not specify a file system on the command line?

Answer:

The command `fsck -p` would search the `/etc/checklist` file for the file systems to be checked. The first field is the device file name of the file system and the sixth field indicates the *pass number* or the order in which to check the file systems. The root file system should have a pass number of 1. All other file systems should have a higher pass number. The significance of this is that file systems on different disk drives can be checked simultaneously if their pass numbers are equivalent in the `/etc/checklist` file. This can save quite a bit of time depending on your configuration.

7. If applicable, run `fsck` on the secondary file system you created and mounted in the module on creating file systems. Use the preening mode of fsck. Note any problems you may encounter and how to fix them.

Answer:

```
# umount /dev/dsk/your_device_file_name
# fsck -p /dev/dsk/your_device_file_name
```

where *your_device_file_name* might be something like `5s0` on a Series 300/400/700 and `c1d0s5` on a Series 800. Of course these names depend on the bus address and section number of the file system you are using.

The description of any errors can be found in the *Error Message Catalog*.

Note that if you used -P the file system probably would not be checked because it was properly unmounted.

8. Properly shutdown your system and run fsck on your root file system disk. Again, note any problems you may encounter.

Answer:

```
# shutdown 0
.
.               (Wait for the system to reach the quiet single-user mode)
.
# fsck /dev/dsk/your_root_file_system_device
```

Remember that the device file name for any mounted file system can be found using either the mount command or the bdf command. On a Series 300/400 this name is most likely /dev/dsk/0s0, on a 700 it is /dev/dsk/6s0 and on an 800 it might be /dev/dsk/c0d0s13.

Module 10 — Backing Up File Systems

Objectives

Upon completion of this module, you will be able to do the following:

- Describe several backup strategies for a system.
- Backup and restore files with `tar` and `fbackup/frecover` using SAM.

About This Module

As a system administrator, backing up your system's files is one of your most important jobs. If a user loses data or accidentally deletes an important file, he or she will come to you and ask you to retrieve the file.

There are many backup strategies and many ways to perform backups and restore files. This module describes the different strategies for backing up your system, and teaches you the commands to backup and restore files.

Overview of Module 10

We just want to give the students enough information so they can effectively do system backups and recover files from a backup. `tar` is presented because of its usefulness in doing small, unscheduled backups. SAM is presented as the way to do system backups. Other backup tools are discussed just because they commonly appear in the manuals and in the UNIX world in general.

10-1. SLIDE: Backup Strategies

Backup Strategies

- Backup the entire file system (full backup)
- Backup part of the file system
  - Files that have changed since the last backup (incremental or delta backups)
  - A subtree of the file system (/users)
- Use a mixed strategy
  - Full backup once a week
  - Incremental backups daily

50722 10-1 99

Student Notes

As mentioned near the beginning of the course, one of the principal responsibilities of a system administrator is preserving the data stored on the system. Unfortunately data is sometimes lost. A piece of hardware may fail, a file may be accidentally removed or overwritten, a command may go astray, or the system may crash. The user community has a reasonable expectation that the administrator has planned and implemented regular backup procedures to minimize the data loss.

Before we can discuss *how* to back up the system, we must first discuss *when* to make backups and *what* to back up. There are two lines of thought concerning what to back up. You can selectively back up parts of the system, or you can back up the entire system. Most of the time, you should implement a backup strategy based on a combination of these two strategies.

As an example, you could do a full system backup (or archive backup) once a week which will store every file on your disk(s) onto the backup media (usually tape). Incremental (or delta) backups can then be done once each day to backup the files that have changed since the full backup was made. You may also decide to backup only a part of the file system tree (like all of the `/users` directory) during the incremental backups.

There are three things to consider when choosing a backup strategy that is right for you:

- amount of media needed
- amount of time needed
- how often backups are needed

The amount of media and the amount of time require that you be able to approximate the amount of space being used by files on the system. You can use the `bdf` command to get the number of free disk blocks. The `bdf` command reports in 1024-byte blocks.

`bdf` might look something like this on a Series 300/400/700:

```
$ bdf
File system    kbytes  used    avail    capacity  Mounted on
/dev/dsk/0s0 487022  142518  295801     33%    /
```

The Series 800 `bdf` will only differ because of the device name for the file system:

```
$ bdf
File system        kbytes  used    avail    capacity   Mounted on
/dev/dsk/c0d0s13 580734  375387  147273     72%     /
/dev/dsk/c3d0s2  580318  173800  348486     33%     /users
```

How often you make backups depends on how much data you can afford to lose. If you can afford to lose a month of data, then you need only back up the system once each month. If you can only afford to lose 6 hours of data, then you must back up every 6 hours. However, backing up every 6 hours can become prohibitive, and other possibilities (such as redundant systems or mirrored disks) must be considered. For most applications, full backups once each week and incremental backups each night are sufficient.

10-1. SLIDE: Backup Strategies

Instructor Notes

Purpose

To introduce the importance of backups and describe some basic strategies.

Key Points

- Stress that each system administrator needs to decide the frequency of backups.

Teaching Tips

- You may want to solicit information from the students about how they currently perform system backups.

Transition

There are many backup tools available on HP-UX. Here are some of the most popular.

10-2. SLIDE: Comparison of Backup Methods

Comparison of Backup Methods

	fbackup	cpio	tar	dd
Backup Type	Logical	Logical	Logical	Physical (image copy)
Speed	Fast	Medium	Slow	Fast
Flexibility	High	Medium	Low	Low
Handles Multiple Media	Yes	Yes	Yes (on HP-UX)	Yes
Portability	HP-UX only	Unix	Unix/some Non-Unix	N/A
Special files	Yes	Yes	No	Yes
Network "smart"	Yes	No	No	No
Can Append Files	No	No	Yes	Yes
Interruptible	Yes	No	No	No
Directory Relative Recovery	Yes	If created with relative paths	If created with relative paths	No
Application	Full and incremental system backups	Copying directory structures and transferring files to other systems	Easy to use and highly portable	Duplication of bit images on to different media

50722 10-2 100

Student Notes

All of these utilities are available on HP-UX. You can pick the utility that best suits your needs. For the purpose of this course we will be discussing the `fbackup` utility as run through SAM for doing your full and incremental system backups and the `tar` utility for archiving a few files or directories just for safe keeping.

Below is a discussion of the most common backup tools:

`fbackup` — `fbackup` and its counterpart `frecover` are the fastest and most flexible tools available for backing up an HP-UX system. These tools can be run using a tape drive that is physically connected to another system on the network. Possibly the best feature of `fbackup` and `frecover`, however, is that they are integrated into SAM. This makes them the easiest tools to use for backups.

`cpio` — A general purpose backup utility in the UNIX world. This was the backup method of choice before `fbackup` because of its flexibility in storing and restoring files and because of its transportability to other UNIX systems.

`tar` — Also a general purpose logical file backup utility, `tar` is the oldest of all the utilities and thus the most portable. It is on all UNIX systems and can be read by several non-UNIX systems as well. `tar` is much easier to use than `cpio` and allows appending files to backup media. `tar` is commonly used for backing up small numbers of files easily.

`dd` — The `dd` command is useful in some limited situations, but is technically *not* a backup command. The `dd` command is a general purpose physical file copy utility. This is different from all the above utilities in that `dd` copies no file names or file attributes to a backup media; it simply copies everything, byte for byte. Thus no selective restore is possible. For these reasons, `dd` is *not* recommended for system backups. It is generally used for two purposes:

- Make a duplicate copy of a disk quickly. This assumes a destination disk the same size or larger than the source.
- Read or translate foreign magnetic tapes. For example, `dd` has the capability for reading backup media with user-defined record sizes, ASCII or EBCDIC translation, byte switching, and other useful options.

Note

Since there are many commands available, and many options to each of these commands, it is strongly recommended that you write the command you use to create your backup on the label of your backup media.

10-2. SLIDE: Comparison of Backup Methods Instructor Notes

Key Points

- Starting with 8.0, the SAM Utility does support backups on all HP-9000 HP-UX systems using the `fbackup` and `frecover` utilities.
- Backing up your file system with `fbackup` is the most flexible method, but it is a proprietary backup not available on non-HP-UX systems.
- Backing up your file system with `cpio` allows the easiest interchange with other Unix systems.
- Using `cpio` has the disadvantage of your not being able to add additional files later, if you attempt this it will destroy the `cpio` archive.
- `tar` spans multiple media as of HP-UX 7.0. `tar` also allows interchange with non-Unix systems.
- `tar` is often considered easier to use than `cpio`.
- There is a new utility designed to work with DDS/DAT tapes called `DATIO`. There is not currently much information on `datio`. You may want to look into this in case it comes up during classroom discussions.

Teaching Tips

Ensure that the students understand the differences between the various utilities and when to use them. Listing out several ways to do the same kinds of thing can sometimes be very frustrating to students. They will probably have to use (at least from a restore standpoint) all of these at one time or another, however, and this is why they are all discussed. We will stick to the easiest methods: SAM for full and incremental backups and `tar` for the rest.

Be warned that the cartridge tape from HP uses a different format than the QIC format used by Apollo Domain or Sun Microsystems SunOS. These formats cannot be used interchangeably in tape drives from the different vendors.

Possibly have the students describe their applications in class and help them to see which utility will work best in their application. Generally they will choose to concentrate on a few rather than all of these backup utilities.

Transition

We will discuss `tar` first for performing a quick backup of just a few files or directories.

10-3. SLIDE: The `tar` Command

The tar Command

Syntax:

`tar` *key* [*file* | *directory* . . .]

Where *key* could be the following:

`f` *device*	Use *device* as the name of the device or file where the archive will be written (`/dev/rmt/0m` is used if the `f` key is not used).
`c`	Create a new archive
`r`	Append to end of existing archive
`x`	Extract from archive
`t`	List contents of archive
`v`	Verbose

50722 10-3 101

Student Notes

`tar` is the easiest tool to use when backing up only a few files or directories at a time for safe keeping. For example, if a user just finished a project and wanted you to make a separate tape to keep and possibly add more files to later, `tar` is the right tool to use. `tar` is also the best way to transfer tape information between different types of non-HP-UX machines.

Its actions are controlled by the *key* argument. The *key* string can be preceded by a hyphen (-) (as when specifying options in other HP-UX commands), but it is not necessary.

If the `f` key is used, it causes `tar` to use the next argument as the name of the archive instead of `/dev/rmt/0m` (which is the default). If the name of the archive is `-`, `tar` writes to the standard output or reads from the standard input, whichever is appropriate. Thus, `tar` can be used as the head or the tail of a pipeline.

Examples:

```
# tar cv file1 file2
# tar tv
# tar rv file3
# tar xv file2
# tar cvf - /users/ralph | tcio -oZ /dev/rct/c3d0s2
```

For a complete list of tar keys and options, see `tar(1)`.

Cartridge Tape Users

The `-f -` option to `tar` is useful if you are writing to a cartridge tape because the output of tar should be piped into the `tcio` command as shown on the slide to write to the cartridge tape drive. Writing directly to the cartridge tape without `tcio` is very slow and can eventually wear out our drive.

10-3. SLIDE: The tar Command — Instructor Notes

Purpose

To explain how to use the tar command.

Key Points

- The tar command was originally designed to archive library files to tape. There is an HP-UX command called ar that archives library files to disk. (tar is a **tape** version of ar, hence its name.)
- If f is specified, it must be followed by an argument that is the name of the device file that you will be using for the archive. If you are writing to the tape drive named /dev/rmt/0m there is no need to use the f option.

Teaching Tips

Go over the examples. They are (in order):

- # tar cv file1 file2

  Creates an archive of file1 and file2 on /dev/rmt/0m. Note that the f key was not given so the device name defaults to /dev/rmt/0m.
- # tar tv Lists the table of contents of the archive
- # tar rv file3 Appends file3 to the archive
- # tar xv file2 Extracts file2 from the archive
- # tar cvf - /users/ralph | tcio -oZ /dev/rct/c3d0s2 Archives the directory /users/ralph to a cartridge tape.

  To read from the cartridge tape: # tcio -i /dev/rct/c3d0s2 | tar *key*f -

  Note that tcio -o write *out* to a tape and tcio -i reads *in* from a tape. The -Z option makes the archive such that you do not need tcio to read the contents.

Transition

While tar is best suited for simple file backups, fbackup run through SAM is the best way to do system backups containing many files.

10-4. SLIDE: Using SAM to Backup and Recover Files

Using SAM to Backup and Recover Files

```
SAM                          Backup and Recovery

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

          Back up Files Interactively ...
          Add an Entry to the Automated Backup Schedule ...
          View/Remove Entries in Automated Backup Schedule ...

          Recover Files or Directories ...
          Get List of Files from a Backup Tape ...

          Show Backup and Recovery History
          View Information on Last:
               Full Backup
               Incremental Backup
               File Recovery
```

Help	Main Menu	Shell	Select Item	hpterm				Previous Menu

50722 10-4 102

Student Notes

The System Administration Manager utility is useful in most backup situations. SAM uses the `fbackup` and the `frecover` commands for all HP-UX systems. These *are* the recommended tools for system backups.

SAM will take you to this screen if you run:

```
# sam (Return)

     Backup and Recovery ->
```

The screen shown on the slide, allows you to select the variety of backup and recovery functions available through SAM.

- `Backup Files Interactively`

  From this screen you can pick and choose the directories to be included in and excluded from the backup. You will watch SAM perform the backup to the specified device and you will not be able to exit SAM until the backup completes.

- `Add an Entry to the Automated Backup Schedule`
- `Remove an Entry from the Automated Backup Schedule`
- `View Currently Scheduled Backups`

  These choices allow you to set up cron jobs so SAM can do your backups on a schedule whether you are there watching or not. You must have a tape in the drive when you set up the backup schedule and whenever the backups are to be run.

- `Recover Files or Directories`

  This will allow you to select files from a backup tape to be restored. There are many options available when restoring files.

- `Get List of Files from a Backup Tape`

  List the table of contents of a backup tape.

- `Show Backup and Recovery History`

  Display the log file that SAM keeps on the backup schedule.

- `View Information on Last: Full Backup`
- `Incremental Backup`
- `File Recovery`

  Displays the index files that SAM keeps which list the files that were backed up or recovered.

10-4. SLIDE: Using SAM to Backup and Recover Files

Instructor Notes

Purpose

To introduce the backup and recovery options available in SAM.

Transition

Let's look at backing up files interactively first.

10-5. SLIDE: SAM — Backup Files Interactively

SAM — Backup Files Interactively

```
SAM                        Back up Files Interactively

   Fill in or modify the desired fields and then press "Perform Task".

        Files to be Included                 Files to be Excluded

   /users/ralph                          /users/ralph/junkdir

   Device file (device to store files on)  /dev/rmt/0m

   Create index file (y or n) . . . . . .  y

 Help | Main | Shell | Perform | hpterm |  |  |  | Exit
      | Menu |       | Task    |        |  |  |  | Task
```

50722 10-5 103

Student Notes

In this screen you must list the directories and files you want to backup, and any sub-directories you want to exclude from the backup. All names must be the absolute pathnames.

You have the opportunity to specify the device file you want to use. This is the device file name of the tape drive you want to backup to. The `Help` screen will tell you what backup devices are available on your system.

SAM also asks you if you want to create an index file. This index file contains list of files that were backed up and which tape it was stored on (if the backup spanned more than one tape). You can view these logs using the `View Status...` menu options. If you want to edit or print these index files they can be found in `/usr/sam/backup/index.full` (from a full backup) or `/usr/sam/backup/index.incr` (for an incremental backup).

If you fill in the form and then select `Perform Task` you will get a confirmation message to continue.

At this point you should be sure that a suitable backup medium is placed in the backup device. It should be ready to use before you continue. You should select `y` to continue the backup process. After you do this, the SAM interface will disappear and you will see the `fbackup` messages as it runs.

In the example on the slide, we chose to backup `/users/ralph` but to leave `/users/ralph/junkdir` out of the archive. The messages we would get from SAM might look like this:

```
Starting interactive backup ...
fbackup(1421): no history is available for graph file
/usr/sam/backup/grapha067 (below level 0)
fbackup(1004): session begins on Thu Jul  9 19:10:06 1992
fbackup(3212): writing volume 1 to the output file /dev/rmt/0m
    1: /
    2: /users
    3: /users/ralph
    4: /users/ralph/.exrc
    5: /users/ralph/.profile
    6: /users/ralph/project
    7: /users/ralph/project/file1
           .
           .
           .
fbackup(1005): run time: 21 seconds
```


Once the backup completes, press the space bar and you will return to the SAM interface.

10-5. SLIDE: SAM — Backup Files Interactively Instructor Notes

Purpose

To introduce the SAM recovery and backup screen menu, and illustrate the easy to learn and flexible support for file system backup that SAM supports.

Key Points

- SAM will allow you to do backups and recoveries across the network if you specify *hostname:device_name* as the backup device. See fbackup(1M) for more information.
- SAM will *NOT* allow you to backup NFS mounted file systems interactively. To do this fbackup requires the `-n` option.
- SAMs automated backups can be made to backup NFS mounted file systems by editing the file `/usr/sam/backup/backauto.scr`. The `-n` option must be added to the `fbackup` command line.
- SAM will not properly backup clustered systems either. Provisions have been made for this, however, in the `backauto.scr` program. You must find the line that looks like this: `# hidden_option="-H"` and remove the `#` from the beginning of the line.

Teaching Tip

- You may want to note what it means for a tape drive to be ready. On a DDS/DAT drive upper light (or left depending on model) must be green and must not be flashing. A yellow light indicates that the tape is write protected (the write protection tab on the back of the tape is open). Mag tapes typically have an on-line indicator that must be on and, of course, you should not see the write protect light on.

Transition

Scheduled backups are a little more difficult because you have to manage when the backup occurs.

10-6. SLIDE: Performing Scheduled Backups

Performing Scheduled Backups

```
SAM              Add an Entry to the Automated Backup Schedule

     Fill in or modify the desired fields and then press "Perform Task".

     Backup type (mark one with an "x") . . . . .      █ full
                                                       _ incremental

     Backup scope (mark one with an "x")  . . . .      _ entire system
                                                       _ selected files

     Choose backup time by (mark one with an "x")      _ days of the week
                                                       _ days of the month

     Device file (defive to store files on) . . .      /dev/rmt/0m________

     Create index file? (y or n)  . . . . . . . .      y_

     Mail results to user (login name)  . . . . .      root__

 Help | Main | Shell | Perform | hpterm |  |  |  | Exit
      | Menu |       | Task    |        |  |  |  | Task
```

50722 10-6 104

Student Notes

Most of the time you will have a schedule for performing your system backups. The screen on the slide shows how SAM can set up your backup schedule so all you need to do is put a tape in the drive at the proper time and SAM will do the backup automatically.

Notice on this screen you can select a `full` or `incremental` backup. Putting an `x` next to `full` will cause SAM to backup every included file and directory. If you select `incremental` instead, SAM will only backup those included files that have changed since the last full backup.

Selecting the `entire system` option will include all files and directories on all of your mounted file systems (NFS file systems excluded) in the backup. If you choose `selected files` you will get a screen so you can list the names of the files and directories to be included in and excluded from the backup.

Each of the days of the week and days of the month options bring up a subwindow that prompts you for the backup time and the weekdays or dates as appropriate. Simply check off the days you would like the backups to occur. This is a sample of the days of the week screen:

```
SAM                              Schedule

     Fi   Enter the time in 24-hour format.
          For example, enter "20:00" for 8 p.m.       ress "Perform Task".

     Ba   Time: 20:30                                  x full
                                                       _ incremental
          Place an "x" beside the correct days.
     Ba   Press "Done" when you are finished.          x entire system
                                                       _ selected files
          Every Day           _
     Ch   Weekdays (M-F)      x                        x days of the week
                                                       _ days of the month
          _ Sunday                x Thursday
     De   x Monday                x Friday             /dev/rmt/m
          x Tuesday               _ Saturday
          x Wednesday
     Cr                                                y

     Mail results to user (login name) . . . . .       root

 Help |     |     | Done | hpterm |     |     |     | Exit
      |     |     |      |        |     |     |     | Window
```

The "Device file" and "Create index file" choices are the same as we have already discussed. Make sure there is a tape in the drive before you press perform task. SAM needs to see the tape so it can properly configure the automatic backup program.

Note that since you may not be present when this backup runs, SAM will mail the status of the backup to the specified user. Typically, this user is root. The specified person can use any of the mailers to read this mail. You can also view the status of *every* backup or recovery using the `Show Backup and Recovery History` option on the "*Backup and Recovery*" screen.

Once you press `Perform Task` your automated backup is submitted to `cron`.

You can later use `View/Remove Entries in Automated Backup Schedule ...` to see the schedule or remove an incorrect entry.

```
SAM                 View/Remove Entries in Automated Backup Schedule

Move the cursor to the line of the entry to be deleted and then press
"Perform Task". Use the tab key to move between fields.

            Day of        Day of       Full/
Time        Month         Week         Incr.      Files (i=included; e=excluded)

20:30  *                  1-5          FULL    i  /
```

Help	Main Menu	Shell	Perform Task	hpterm				Exit Task

10-6. SLIDE: Performing Scheduled Backups — Instructor Notes

Purpose

This is probably the most *important* part of this module.

Key Points

- The fields are pretty self-explanatory.
- There must be a tape in the specified device in order to correctly complete this screen. If you are backing up to DAT, for example, and there is no tape in the drive, SAM will assume that you are backing up to a cartridge tape and configure the automatic backup program to use `tcio`. This will cause the backup to fail when it runs.

Teaching Tips

- These screens are more easily shown than described. Walk the students through the screens so they can see the different fields.

Transition

Now that we have a backup made, we would like to be able to see what was backed up.

10-7. SLIDE: Viewing A Backup Index

Viewing A Backup Index

```
SAM                              Backup and Recovery

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

          Back up Files Interactively ...
          Add an Entry to the Automated Backup Schedule ...
          View/Remove Entries in Automated Backup Schedule ...

          Recover Files or Directories ...
          Get List of Files from a Backup Tape ...

          Show Backup and Recovery History
          View Information on Last:
                  Full Backup
                  Incremental Backup
                  File Recovery
```

Help	Main Menu	Shell	Select Item	bpterm				Previous Menu

50722 10-7 105

Student Notes

The highlighted choices on the slide are the different ways you can view what was backed up (or recovered).

Whenever you make a backup with SAM it will create an index file at the beginning of the tape. This index can be retrieved at any time by loading that tape and selecting `Get List of Files from a Backup Tape ...`. SAM prompts you for the device name and the name of a file to place the tape's index in. Beware that an index file can very easily be larger than 1Mb for a complete system backup. When SAM reads the index from tape it is displayed on your screen and saved into the specified file.

The index file specifies the tape number and the pathname of the files and directories that were backed up. An index file will usually look something like this:

```
#  1 /
     .
     .
#  1 /users
#  1 /users/ralph
#  1 /users/ralph/.kshrc
#  1 /users/ralph/.profile
#  1 /users/ralph/project
#  1 /users/ralph/project/dir1
#  1 /users/ralph/project/dir1/fileX
#  1 /users/ralph/project/file1
#  1 /users/ralph/project/file2
     .
     .
```

SAM will also ask you if you want to create an index file when you perform a backup. If you answer y to this question, the list of files backed up is kept in a file *on the disk* in addition to the list written on tape.

Note that this disk index file is overwritten each time you do another backup. The index files are in the directory `/usr/sam/backup`. Their names are `index.full` and `index.incremental`. You can copy these files to another location or print them for safe keeping.

The index file is *always* written to tape even if you answer n to the index file question when you make the backup.

When you select `View Information on Last: File Recovery` you will see a listing of all files that were restored to the disk from the tape:

```
      .
      .
drwxr-xr-x    ralph   users   /users/ralph/project
drwxr-xr-x    ralph   users   /users/ralph/project/dir1
-rw-r--r--    ralph   users   /users/ralph/project/dir1/fileX
-rw-r--r--    ralph   users   /users/ralph/project/file1
-rw-r--r--    ralph   users   /users/ralph/project/file2
      .
      .
```

10-7. SLIDE: Viewing A Backup Index — Instructor Notes

Purpose

This is very useful for finding which files were backed up on which tape. You also need to know the exact pathname of the files if you want to recover only certain files from the tape. This is a good way to get that pathname.

Key Points

- The index file is always written to tape.
- The index file on disk is more accurate than the one on tape because the tape index file is written *before* anything is actually backed up. In other words, it is a list of what *should* be on the tape. The optional disk index file is a list of what really *is* on the tape. Comparing the disk index file and tape index file is a good way to see if everything was really backed up.
- If a backup spans multiple tapes, the index file on the last tape is the only one that contains the correct volume numbers on the files.

Transition

We know how to make a backup tape and list its contents, but how can we retrieve files from a backup tape?

10-8. SLIDE: Recovering Files

Recovering Files

```
SAM                          Recover Files or Directories

   Fill in or modify the desired fields and then press "Perform Task".

   Specify files to recover by (mark one with an "x"):

                    x entering names of individual files or directories
                    _ choosing a file containing a list of files

   Overwrite newer files? (y or n) . . . . . . . . . .  n

   Maintain original file ownership? (y or n)  . . .  y

   Recover files using full path name? (y or n)  . .  y

   Place recovered files relative to this directory:  /

   Device file (device to get files from)  . . . . .  /dev/rmt/0m
```

Help	Main Menu	Shell	Perform Task	hpterm				Exit Task

50722 10-8 106

Student Notes

Doing backups faithfully means nothing if you do not know how to recover the backed up files. SAM can help us do this using the screen shown on the slide.

```
# sam (Return)

      Backup and Recovery ->

            Recover Files or Directories ...
```

The easiest way to recover something from a backup tape is by "`entering names of individual files or directories`". When you select this item, SAM brings up a screen that allows you to name the files and directories to recover. You must specify the pathnames exactly as you did when the tape was made (this should be the absolute pathname). You may want to view the index of the backup tape first to see the actual pathnames on tape.

Normally older files from a backup tape will **not** overwrite the same disk file that may have changed. If you suspect that a file or directory is somehow corrupt, then you should answer y to the "`Overwrite newer files?`" question.

"`Maintain original file ownership?`" is the preferred way to recover most files. If you answer "n" to this question, `root` will own all of the recovered files.

The next two questions affect where the files will be restored. The default answers will recover selected files into the same directory they were in when the backup was made.

If you answer "y" to the question "`Recover files using full path name?`" the files and directories will be restored into the directory specified in the next field (/ by default). The "tree" of files and directories will be rooted in this directory and it will maintain its shape. Most of the time this is the type of recovery you want.

Answering n to the full path name question will recover all of the specified files into the specified directory without creating any subdirectories. *All* of the specified files are in *one* directory. This could be desirable if you want to restore a project directory into a different directory than the one it came from when it was backed up.

Example Recovery

Here is an example of a directory recovery to help clarify where files are restored onto the disk.

Assume you made a full system backup of the entire file system. You want to restore a directory structure from the backup tape. The directory was backed up as `/users/ralph/project`. The directory tree originally looked like this:

```
        /users
          |
         ralph
          |
        project
    -------|-------
    |      |      |
  file1  file2  dir1
                  |
                fileX
```

To recover this directory structure you must specify `/users/ralph/project` as the name of the directory to be recovered (put this pathname in the `include` column).

Now you have three choices regarding where the directory and its files are recovered:

1. If you want the directory and its contents restored to `/users/ralph/project` (their original location) you would answer y to the full pathname question and specify / as the directory to place the files in. The tree that was recovered would look exactly like the original. This is the simplest and most common type of file recovery.
2. What if you wanted the entire directory tree restored into a different directory like `/tmp`? You would answer y to the full pathname question and specify the directory name as `/tmp`. The restored files would look like this:

```
            /tmp
              |
            users
              |
            ralph
              |
           project
       -------|-------
       |      |      |
     file1  file2  dir1
                     |
                   fileX
```

   Notice that because you wanted the recovery to complete using full pathnames the resulting directory tree is identical to the original except that it is rooted in the specified directory (`/tmp` in this example).
3. The third option you have is to just recover specified files and not the directories and subdirectories. If you wanted to restore `file1`, `file2`, and `fileX` into norton's home directory you would answer n to the full pathname question and specify `/users/norton` as the directory to place the files in. The resulting tree would look like this:

```
            /users
              |
            norton
       -------|----------- ...   ------------
       |      |      |              |   |   |
     file1  file2  fileX   ...  norton's other files
```

   Notice how the `project/dir1` directory was not recovered but `fileX` was. When you answer n to the full pathname question *all* of the specified files are recovered into the specified directory and no subdirectories are created.

10-8. SLIDE: Recovering Files

Instructor Notes

Purpose

Students must be familiar with recovering files.

Key Points

- You cannot specify relative pathnames when backing up files with SAM. Therefore, the choices for recovery are restricted to those demonstrated in the student notes.

Teaching Tips

- Stress that without a good knowledge of recovery, backups have little use.
- Go through the three different types of recoveries. You might do this or other examples on the board to help clarify the choices you have.

Transition

Now for some practical experience with these backup commands.

10-9. WORKSESSION: Review Questions

1. What type of archive is the `tar` command used for?

2. What is the main function of the `fbackup` command?

3. If you do not specify the `f` key to `tar`, where is the archive written?

4. Can you append files or directories to an `fbackup` archive tape? What about a `tar` tape?

10-9. WORKSESSION: Review Questions

Instructor Notes

1. What type of archive is the `tar` command used for?

Answer:

Simple backups of a few files or directories. It is too slow to use it as an effective system backup tool.

2. What is the main function of the `fbackup` command?

Answer:

To perform full and incremental system backups. It is very fast compared to other tools available on the system.

3. If you do not specify the `f` key to `tar`, where is the archive written?

Answer:

To the default device `/dev/rmt/0m`.

4. Can you append files or directories to an `fbackup` archive tape? What about a `tar` tape?

Answer:

You cannot append files to an `fbackup` archive.

Appending files to an archive is one of `tar`'s best features.

10-10. LAB: Backing Up the System

Directions

Ask your instructor about the available devices in your classroom for performing these labs.

Add a user named norton to your system using SAM. We will use norton's home directory as an example for the backup labs. You may also want to make directories and copy some extra files into norton's home directory to make the exercises even more interesting.

1. Using the `tar` command, backup norton's home directory.

2. List the contents of the `tar` tape.

3. Append the file `/etc/passwd` to the tape.

4. Remove norton's `.profile` from his home directory then Recover *just* norton's `.profile` from the tape.

5. Interactively backup just the `/users` directory using SAM. Make sure you create a *disk* index file.

6. View the index file from the tape then view the index file from the disk. Are they the same?

7. Make a new directory called `/users/testdir` then recover just norton's home directory into `/users/testdir`.

10-10. LAB: Backing Up the System — Instructor Notes

Directions

The `tar` exercises can be done using a regular file name as the device instead of a tape device.

SAM will not accept a regular file as the backup device. For this reason you will have to have the students take turns performing backups on the available devices in the classroom for the SAM exercises.

If you are running short of time and/or you do not have a reasonable number of tape devices, you could run a demonstration of a backup and recovery using SAM. Students should really try this, however, because it is so important.

1. Using the `tar` command, backup norton's home directory.

Answer:

```
# tar cvf device_name /users/norton
```

OR better:

```
# cd /users/norton
# tar cvf device_name .
```

device_name will be the name of your tape device. If your tape device is named `/dev/rmt/0m` you do not need the `f` *device_name* on the command line.

If there are not enough tape drives for every system in the class, you can specify a regular filename for the *device_name*.

By cd'ing to the directory to backup first then archiving "`.`", you can recover the files into any directory. Not just `/users/norton`.

2. List the contents of the `tar` tape.

Answer:

```
# tar tvf device_name
```

3. Append the file `/etc/passwd` to the tape.

Answer:

```
# tar rvf device_name /etc/passwd
```

4. Remove norton's .profile from his home directory then Recover *just* norton's .profile from the tape.

Answer:

```
# tar xvf device_name /users/norton/.profile
```

If you created the archive using "." instead of /users/norton you could go to any directory on the system and recover the .profile into that directory:

```
# cd /users/norton
# tar xvf device_name .profile
```

5. Interactively backup just the /users directory using SAM. Make sure you create a *disk* index file.

Answer:

On the Back Up Files Interactively... screen you must specify /users to be included in the backup. You could exclude someone's home directory from the backup by putting /users/*user_name* in the exclude column if you wanted to experiment more.

Answer y to the "Create index file?" question.

6. View the index file from the tape then view the index file from the disk. Are they the same?

Answer:

To get the index file from tape you must use **Get List of Files from a Backup Tape ...** and specify the correct device name. The index will be displayed on your screen and saved in the file you specified.

To see the index file created on disk during the backup select
View Information on Last: Full Backup

7. Make a new directory called /users/testdir then recover just norton's home directory into /users/testdir.

Answer:

On the "*Recover Files or Directories*" screen you need to "enter names of individual files or directories" and specify /users/norton in the include column.

The directory for recovery should be /users/testdir.

If you answered y to the full pathname question, you would recover /users/norton into /users/testdir. So you would actually see /users/testdir/users/norton/*norton's_files* in a listing.

Answering n to the full pathname question, would recover just *norton's_files* directly into /users/testdir.

Module 11 — Other Administration Tasks

Objectives

Upon completion of this module, you will be able to discuss or explain the following:

- Scheduling processes to run with `cron` and `crontab`
- Reconfiguring the kernel with `sam`
- The `mkrs` command for the Series 300/400/700
- The Series 700 restore tape

About This Module

This module focuses on some of the other tasks that you may need to do as part of you job as system administrator. None of the topics presented here is large enough to have its own module so they are combined.

Do not discount the importance of these tasks because they are combined into the *last* module of the course. On the contrary, tools like `cron` are commonly used to help you manage your system even if you are not there. The recovery systems could be the difference between fixing a small problem on your root file system disk and reinstalling it to fix the problem.

Overview of Module 11

These topics did not fit anywhere else but are useful (or necessary) things to do. `Cron` is generally applicable to everyone. Reconfiguring the kernel is something that most students will not have to do but if they should have to reconfigure someday for an application, this module shows how SAM can be used to do that. The recovery mechanisms are only for the workstation users.

11-1. SLIDE: Scheduling Programs Using cron

Scheduling Programs Using cron

1. Root must add login name to /usr/lib/cron/cron.allow

```
# vi /usr/lib/cron/cron.allow
```

2. User creates cronfile containing programs to be submitted

```
$ vi cronfile
```

3. User submits cronfile to cron process with crontab command

```
$ crontab cronfile
```

50722 11-1 107

Student Notes

Regular users, as well as root, can utilize cron for executing programs on a schedule. For example, you might use cron to schedule your daily backup routines (this is how SAM schedules backups). The jobs will execute whether or not you are logged in. Any output from your jobs will be electronically mailed to you.

The cron process is normally started during boot-up by the /etc/rc script. Therefore, as long as you are in a multi-user run-level, cron should be executing on your system. If for some reason it is not, start it running by typing:

```
# cron
```

Jobs are submitted to cron with the crontab command. Root controls who can use crontab through the /usr/lib/cron/cron.allow file. Users are permitted to use the crontab command if their names appear in the cron.allow file. If this file does not exist or is empty then only root is allowed to use the cron facility.

Once you have access to the crontab command, you must create a file that contains the commands you wish to submit to cron and the times you want them to be executed. We will look at the format of this file on the next slide. By convention this file is called cronfile but you can call it anything you want.

To submit your job to cron, type:

```
$ crontab cronfile
```

Once you submit your file with the crontab command, your job is activated on the specified schedule. If you run crontab again using a different file the original jobs will be replaced with the new ones.

To see the jobs you have scheduled with cron use the crontab -l command.

To remove jobs from cron use the crontab -r command.

11-1. SLIDE: Scheduling Programs Using cron — Instructor Notes

Purpose

To give an overview of how to submit and control jobs with the cron system.

Key Points

- cron should be running on the system if it is in multi-user mode.
- If cron.allow does not exist, then /usr/lib/cron/cron.deny is checked to determine if the user should be denied access. If both exist, cron.allow takes precedence. If neither file exists, only root is allowed to submit a job. An empty cron.deny file allows all users to use crontab.
- Each user may have only one cronfile because running crontab destroys all previously submitted jobs.

Files That Grow Without Bound

There are many files that grow with out bound, including /etc/wtmp and /etc/btmp. Files such as these should be monitored by the system administrator and appropriate action should be taken on a regular basis. You could do this with the cron facility. Below is a list (not necessarily complete) of files that tend to grow without bound.

Table 11-1.

File Name	Written to by	Read by
/etc/wtmp	/etc/login	/etc/last
	/etc/init	
	/usr/lib/acct/accton	/usr/lib/acct/accton
/etc/btmp	/etc/login	/etc/lastb
/usr/adm/sulog	/bin/su	
/usr/lib/cron/log	/etc/cron	
/usr/spool/mqueue/syslog	/usr/lib/sendmail	/usr/bin/mailq
/usr/spool/mqueue/sendmail.st		
/usr/spool/uucp/diallog	Many of the UUCP	
/usr/spool/uucp/errlog	commands such as	
/usr/spool/uucp/logfile	/usr/lib/uucp/uuxqt	
/usr/spool/uucp/syslog	/usr/lib/uucp/uucico	
/usr/spool/uucp/culog		
/usr/spool/lp/log	/usr/lib/lpsched	/usr/bin/lpstat
	/usr/bin/lp	
/usr/adm/messages	/etc/dmesg	

These files could be trimmed by a weekly cron job. An example cronfile to trim `/etc/wtmp` every week might look like this:

```
23 59 * * 0    cp /etc/wtmp /etc/OLDwtmp; > /etc/wtmp
```

The "`> /etc/wtmp`" command will truncate the `wtmp` file to 0 bytes. This is better than removing the file and creating it with `touch` because you don't change the file permissions or ownership.

Transition

Let's take a look at the format for `cronfile`.

11-2. SLIDE: cronfile Format

cronfile Format

- Consists of lines of six fields each. The first five fields specify the following:
  - minutes after the hour (0-59)
  - hours of the day (0-23)
  - days of the month (1-31)
  - months of the year (1-12)
  - days of the week (0-6 with 0=Sunday)
- An "*" indicates all legal values, "-" is a range and "," includes several values.
- Fields are separated by spaces or tabs.
- The last field is the program to be executed.
- Example:

```
0 * * * *          /bin/date  >  /dev/console
0 6 * * 1,3,5      /etc/lastb | lp
30 7 1-7 * 2       monthly_report
30 20 * * 0        weekly_full_backup  >  /tmp/backup.log
```

50722 11-2 108

Student Notes

The entries in `cronfile` must be in a specific format to be interpreted successfully by `cron`. Each entry in the file is a line containing six fields separated by spaces or tabs.

The first five fields contain integers which represent the date and time a command is to be executed. They are shown on the slide. Each of these fields may contain an asterisk which represents all legal values, or a list of numbers separated by commas to include all of the listed values or the two numbers may be separated by a dash which specifies a range.

The last field is the command line that is executed by the shell at the specified times.

Some examples of using cron are given on the slide.

Note

If you do not redirect the output and errors of your cron jobs, any output generated will be electronically mailed to you.

11-2. SLIDE: `cronfile` Format — Instructor Notes

Purpose

To show the students how to create a `crontab` file and show some examples of using cron.

Key Points

- Standard output and error should be redirected in the commands or the output will be mailed to the user.
- Your cron file does not have to be named `cronfile`. You can name it whatever you want. `cronfile` is the standard name.

Teaching Tips

- Explain the examples on the slide:

1. writes the date to the console every hour on the hour.
2. prints a list of bad login attempts at 6am, Monday, Wednesday, and Friday.
3. runs `monthly_report` on the first Tuesday of the month.
4. runs the `weekly_full_backup` command at 8:30pm every Sunday.

Transition

What if we want to change the jobs we have submitted?

11-3. SLIDE: Modifying Your cronfile

Modifying Your cronfile

1. List contents of existing `cronfile`. (First move to the directory where your `cronfile` resides.)

```
$ crontab -l > cronfile
```

2. Make changes to `cronfile`.

```
$ vi cronfile
```

3. Re-submit your `cronfile` to `cron`.

```
$ crontab cronfile
```

50722 11-3 109

Student Notes

Once a user's **cronfile** is activated with the **crontab** command, any future use of this command will cause the original crontab entries to be replaced with the contents of the new cronfile.

To change information in your crontab file, you must first retrieve the existing crontab file, modify it, and then re-submit it with **crontab**.

Recall that you can remove all of your jobs using the **crontab -r** command.

11-3. SLIDE: Modifying Your `cronfile` — Instructor Notes

Purpose

To teach the students how to modify their `cronfile` after submitting it to `cron`.

Transition

That is all there is to using cron. Let's move to the next topic: Kernel reconfiguration.

11-4. SLIDE: Reconfiguring the Kernel with SAM

Reconfiguring the Kernel with SAM

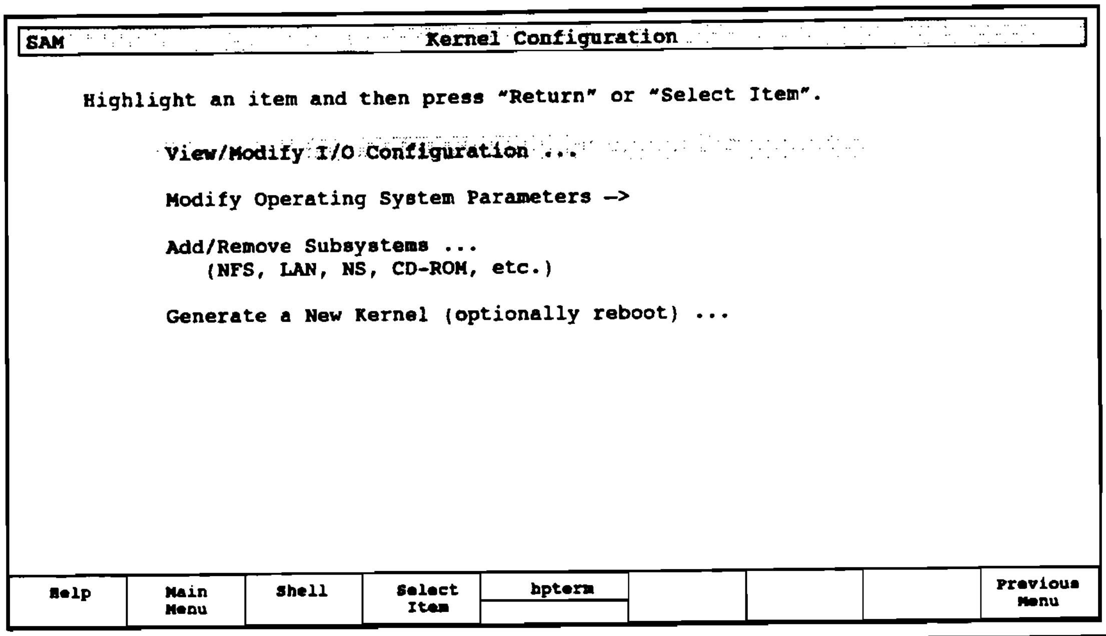

50722 11-4 110

Student Notes

Most of the time the HP-UX kernel will be correctly configured to properly run your applications. There may be times, however, when you add a new application that is larger than the kernel will allow when it runs. You may also be adding a new type of device on your system that requires a new driver be configured into the kernel so it can recognize the device. You may also add a new software subsystem (like LAN or NFS) to your system which would also require a kernel modification.

This slide shows you the kernel reconfiguration tasks that SAM can perform for you. As you can see, SAM is quite versatile in this area. In most cases where you need to reconfigure the kernel, you will probably want to use SAM.

Note

Reconfiguring the kernel requires that the system be rebooted to activate the new configuration. For this reason it is best to shutdown your system to the quiet single-user mode before attempting to reconfigure the kernel.

To get to the menu shown on the slide:

```
# sam (Return)

      Kernel Configuration ->
```

On Series 300/400/700 systems before SAM continues, you will get a message that says:

```
The file below provides current kernel values for viewing or
modification.  To use a different set of initial values,
specify your own file.  In either case, press "Return" or
"Done" to proceed.

File: /hp-ux
```

The file /hp-ux contains your executable kernel. This is the file that SAM will rebuild. You will typically use /hp-ux because you are modifying the values in that file.

11-4. SLIDE: Reconfiguring the Kernel with SAM

Instructor Notes

Purpose

To describe three of the most common reasons for modifying the kernel and to show the functionality provided in SAM for tasks related to reconfiguring the kernel.

Key Points

- SAM can perform all the types of kernel reconfiguration that we would be concerned with here:

  Adding/Removing subsystems
  Adding device drivers
  Changing kernel parameters

- The Series 800 has more choices on this screen than are pictured you may or may not want to mention these:

  Change System Console Type
  Modify Primary Swap
  Disk Mirror Configuration

Teaching Tips

- For each of the types of changes we are dealing with, show the students the submenu to use to make those changes:

  To add or remove a subsystem like NFS you would select `Add/Remove Subsystems...`

  To make changes to run larger applications you would select

  `Modify Operating System Parameters....` Common things to change would include `maxtsiz`, `maxdsiz` and `maxusers`.

  To add a new type of device you would select `View/Modify I/O Configuration....`

Transition

We will look at adding and removing subsystems from the kernel first.

11-5. SLIDE: Adding or Removing a Subsystem with SAM

Adding or Removing a Subsystem with SAM

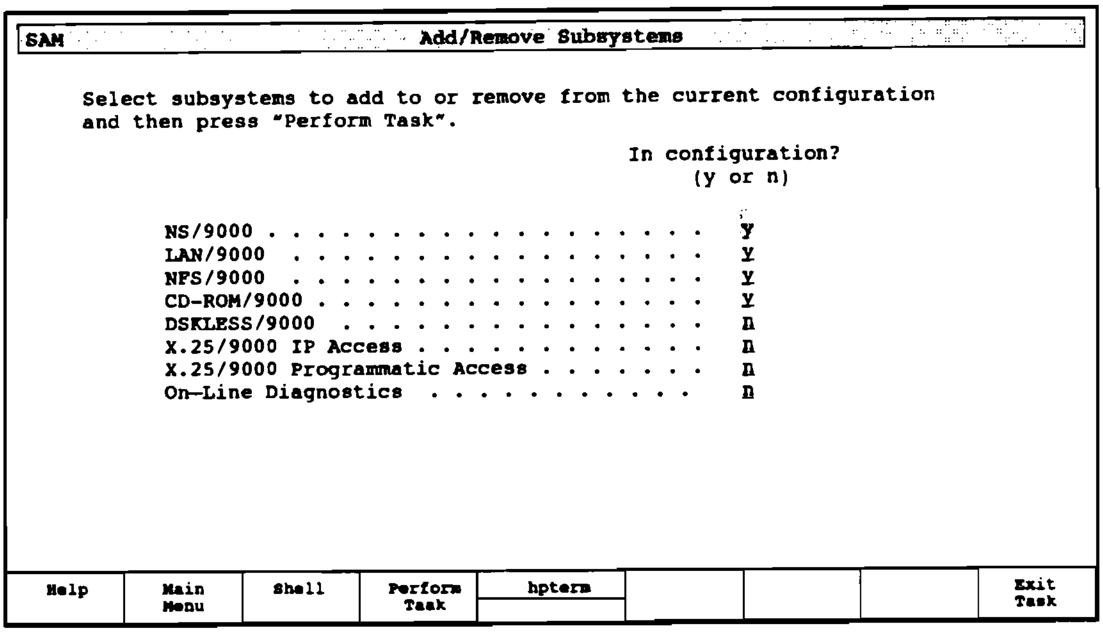

50722 11-5 111

Student Notes

To add or remove a software subsystem like LAN, NFS, or CD-ROM support:

```
# sam (Return)

     Kernel Configuration ->

          Add/Remove Subsystems ...
```

With the arrow keys, move to the subsystem you want to add or remove. Type `y` to include that subsystem in the kernel or `n` and then press `Perform Task`.

You will then get a message similar to the following:

```
Subsystem information has been modified.  You will be
prompted to generate a new kernel to effect changes.

   -- Press the space bar to continue. --
```

Press the Space bar to continue and then select `Generate a New Kernel (optionally reboot) ...` which will tell SAM you have made all the changes you want. SAM will rebuild the kernel and you will be prompted to reboot the system to activate the new kernel.

If you choose not to reboot SAM will give you directions on how to activate the new kernel at a later time.

11-5. SLIDE: Adding or Removing a Subsystem with SAM

Instructor Notes

Purpose

To explain how to add or remove a subsystem in SAM.

Key Points

- Pressing `Perform Task` on the subscreen only tells SAM that you want to make the changes as listed.

  The changes are not made until you select `Generate a New Kernel`. If you try to leave the Kernel Configuration section of SAM without generating a new kernel it will stop you and warn you that your changes could be lost.
- Typically if you load a subsystem using `update` or during the installation process, the kernel will be modified to support this subsystem for you.

Teaching Tips

- Give some examples for adding or removing a subsystem:

  If you just purchased a new CD drive, you will have to add the CD-ROM subsystem so you can use CD based file systems (like the Update or LaserROM disks).
  If you decide you no longer need NFS because you are not sharing file systems with other machines you may want to remove it to reduce the size of your kernel (both on disk and in memory). If you remove software using `rmfn` it will tell you if a kernel change is necessary.

Transition

Operating system parameters are another part of the kernel you may need to change.

11-6. SLIDE: Changing System Parameters With SAM

Changing System Parameters With SAM

```
SAM                 Modify Operating System Parameters

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

                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 ...

Help | Main Menu | Shell | Select Item | hpterm | | | | Previous Menu
```

50722 11-6 112

Student Notes

To get to this menu:

```
# sam (Return)

      Kernel Configuration ->

            Modify Operating System Parameters ...
```

Notice how all of the operating system parameters are divided into groups. To change a particular parameter, you must know what group it is in. You can use SAMs `Help` softkey to help you find what you need to change. Once you know what group your parameter is in, highlight it and press `Select Item`.

If you were to choose `Process Parameters ...`, you would get the following new menu:

```
SAM                         Process Parameters

      Fill in or modify the desired fields and then press "Perform Task".

   Maximum number of processes which may
        simultaneously exist (nproc) . . . . . . . . . . . . 316

   Maximum number of processes that a
        user may have (maxuprc)  . . . . . . . . . . . . . . 50

   Maximum process data segment size in bytes (maxdsiz). . . 67108864

   Maximum process stack size in bytes (maxssiz) . . . . . . 8388608

   Maximum process shared text segment
        size in bytes (maxtsiz). . . . . . . . . . . . . . . 67108864
```

Help	Main Menu	Shell	Perform Task	hpterm				Exit Task

Use the arrow keys to get to the parameter you wish to change and then type in the new number. If you are not sure of what a parameter does or what the appropriate values are, use the `Help` softkey. Press `Perform Task` when you are finished.

Note

You should never modify operating system parameters unless you know exactly what should be changed.

If an application requires a parameter change it will be documented with the notes for installing the application.

If you are getting system error messages, look them up in the *Error Message Catalog*. If you do not know what to change to resolve the problem, call the HP Response Center.

Reference

Appendix A of the *HP-UX System Administration Tasks Manual* describes each tunable kernel parameter in detail.

11-6. SLIDE: Changing System Parameters With SAM

Instructor Notes

Purpose

To show how kernel parameters might be changed if you know what you are doing.

Key points

- Encourage students to carefully read Appendix A of *HP-UX System Administration Tasks* before even considering a change.
- Repeat the warning that you should never modify operating system parameters unless you know exactly what will happen when you make a particular change.
- Don't go over each system parameter in detail, just give the students an idea of what's in there.

Transition

We should also mention how to add new device drivers.

11-7. SLIDE: Adding/Removing I/O Drivers

Adding/Removing I/O Drivers

```
SAM                          I/O Drivers

The I/O drivers currently configured on your system are listed below.
Select the drivers to add or remove and then press "Perform Task".

                                                  In configuration?
                                                      (y or n)
Disk and Tape Drivers:
   SCSI disks (scsi) .................................... y
   SCSI tapes (scsitape)................................. y
   CS80 tapes (cs80)..................................... n

Device Drivers:
   Plotters and device I/O library (hpib)................ n
   Parallel printer/plotter (CharDrv).................... y

System Bus Driver:
   EISA I/O Bus (eisa)................................... y

 Help | Main | Shell | Perform | hpterm |  |  |  | Exit
      | Menu |       | Task    |        |  |  |  | Task
```

50722 11-7 113

Student Notes

To add support for a new type of peripheral to your system, you will have to make sure the appropriate driver for that peripheral is in the kernel. This is done by selecting:

```
# sam (Return)

      Kernel Configuration ->

            View/Modify I/O Configuration ...
```

The screen on the slide is for a Series 700. The other systems will vary only slightly. You simply have to mark a y next to the types of devices you want to support on your system then press `Perform Task`.

Once again, you will have to `Generate a New Kernel (optionally reboot) ...` to actually make the change to the kernel.

11-7. SLIDE: Adding/Removing I/O Drivers — Instructor Notes

Purpose

To show a quick sample of how to add a new device type to the system.

Key Points

- The Series 800 actually has one more menu that allows you to pick the type of device to add (disks, tapes, printers, etc).

Transition

The last topic we want to cover in this module is the recovery system for the workstation family.

11-8. SLIDE: The Series 300/400/700 Recovery System

The Series 300/400/700 Recovery System

A recovery system is useful if:

- The root password is unknown (or /etc/passwd was corrupt or accidently deleted).
- You have a bad boot area on the system disk and cannot boot up.
- There is no login prompt on any device.

You can make a recovery system on:

150 foot Cartridge Tape

SCSI DDS/DAT Tape

Optical Disk

50722 11-8 114

Student Notes

A recovery system is a bootable HP-UX system on some media other than your hard disk. It can be used to boot your system so you can repair some simple problems with your root file system. For example, if you forget your root password there is no way to log into the system as root to change it. Because of this no administration can be done on the system. The recovery system will allow you to remove the root password.

The recovery system can only fix simple problems like a forgotten root password, corrupt superblock, bad boot area or missing system files and programs. If the root file system is unfixable with the recovery system, you will have to reinstall the HP-UX software from scratch and reload the most recent backup(s).

As you can see from the slide you can make your recovery system on a variety of devices. If you choose to make your recovery system on cartridge tape, use a 150-foot cartridge tape. You could use a 600-foot tape, but it takes much longer.

Once you have made the recovery system, *lock it up.* It could be used to break into your system.

You should remake the recovery system every time you install, update to a new revision of HP-UX, or change the kernel configuration of your system.

To make a recovery system, use the `mkrs` command:

`mkrs [-v] [-q] [-s] [-f` *rcdev*`] [-r` *rootdev*`]`

Where:

`-v`	is the verbose option. This is recommended, to see the output of the `mkrs` command as it builds the recovery system.
`-q`	is the quick option. If enough free disk space is available in `/usr/tmp` (typically 10-20Mb), the `-q` option can be used to make `mkrs` create an image of the recovery system in this directory before copying it to the recovery media. This option generally saves a great deal of time when making a recovery system on a DDS/DAT or cartridge tape.
`-s`	is the "small" option. This is used to minimize the size of the recovery system when it is booted. It should be used if your system contains less than 8Mb of memory. Using the `-s` option will make the recovery much more difficult because the tools to perform the recovery tasks for you will not be there.
`-f` *rcdev*	Specifies the device you want to make the recovery system on. Replace *rcdev* with the character device file. If you do not use the `-f` option, `mkrs` assumes you want to use the device `/dev/update.src`. Which is the name of the device you used to initially load the system software.
`-r` *rootdev*	Specifies the device name of your root file system. You can execute the `mount` or `bdf` commands to find the device name of the root file system. If you do not use the `-r` option, `mkrs` will use `/dev/dsk/0s0`.

Series 700 Notes:

If you are building a Series 700 recovery system on DAT the `-s` option is required. However, you will still get the full set of recovery tools when your recovery system is booted.

The `-r` *rootdev* option must always be used on the Series 700. The root file system device name is normally `/dev/dsk/6s0` (not `/dev/dsk/0s0`). If you do not use the `-r` option to `mkrs` the recovery system will not be able to perform the recovery properly.

Examples

The Series 700 `mkrs` command line will normally look like this if you are creating a recovery system on a DAT tape:

```
# mkrs -sqv -f /dev/rmt/0m -r /dev/dsk/6s0
```

The Series 300/400 `mkrs` command line will normally look like this if you are building the recovery system on the DAT drive you used to install your system software:

```
# mkrs -qv
```

Reference

The *HP-UX System Administration Tasks* has the details of building and using the recovery system.

11-8. SLIDE: The Series 300/400/700 Recovery System

Instructor Notes

Purpose

To explain why a recovery system is valuable and how to make a recovery system.

Key Points

- The recovery system can only fix simple problems.
- A `mkrs` recovery system is the source of a potential security breach for your system. It must be keep out of the hands of unauthorized users. Lock it up, possibly with your **UPDATE** and **INSTALL** media.
- The recovery system must be remade when the kernel changes.

Transition

Using the recovery system is much more involved than making it.

11-9. SLIDE: Using A Series 300/400/700 Recovery System

Using A Series 300/400/700 Recovery System

1. If the system is running, execute `sync`, then power off.
2. Put your recovery system on line and make ready.
3. Boot in attended mode and select the recovery system device.

Wait for root file system to be checked for inconsistencies.

When the following menu appears, choose the appropriate option:

```
Select one of the following options by number:

1) Remove the root password
2) Work in a shell to perform recovery manually
3) Perform an automatic recovery
4) Exit recovery system and reboot root file system
5) Help
Selection >> 1
```

50722 11-9 115

Student Notes

Regardless of the problem with your root file system, the procedure for using the recovery system is roughly the same.

Once the recovery system is up and running you will see the menu shown on the slide. (You will not get this menu if you used the `-s` option when creating a Series 300/400 recovery system).

Here are some guidelines for using the recovery system options:

- `Remove the root password`—This is the most obvious and least serious of the choices. The root password is simply deleted from the `/etc/passwd` file on the root file system. Once this task is completed you should `Exit recovery system and reboot root file system` to login as normal.
- `Work in a shell to perform recovery manually`—This requires the most experience from the administrator. It simply starts a superuser shell and you have to find and fix the problem yourself.

- `Perform an automatic recovery`—This will fix almost anything that could be wrong with your system. The boot area of the root disk is replaced and all critical configuration files and programs are copied from the recovery system onto the root file system. The old files are kept in a directory on the root file system called `/tmp/recovery`*something*. The actual name of the directory and all of the files stored there will be displayed on the screen as the recovery system works. You should write down what files are being replaced so you can later check them and put them back if necessary.
- `Help`—This will display a help screen describing the different menu choices.

Reference

Chapter 7: "Backing Up and Restoring Your Data" from *HP-UX System Administration Tasks* explains how to use the recovery tape to restore your system.

If you have any trouble using your recovery system you should immediately contact the HP Response Center.

11-9. SLIDE: Using A Series 300/400/700 Recovery System

Instructor Notes

Purpose

To show the basics of using a recovery system. The process is relatively simple to explain, but time consuming and difficult to actually carry out.

Key Points

- The automatic recovery is not as simple or complete as it sounds.
- In the automatic recovery section, key programs like /bin/sh and /etc/init are replaced. Configuration files like /etc/passwd are also replaced. When you reboot from the root file system you will have to look through the `/tmp/recovery*` directory and compare the original files there with those the recovery system placed on your disk.

  One specific file that will cause serious problems is the `/etc/init` program that comes from the recovery system. When you reboot from your disk, `init` will tell you it cannot read `/etc/inittab`. The problem is that the new `init` program has a problem and you *must* replace it with the original `/etc/init` file.

Transition

Another mechanism to rebuild a problem root disk is the Series 700 restore tape.

11-10. SLIDE: Series 700 Restore Tape

Series 700 Restore Tape

The `restore` operation is provided as a recovery mechanism in the event that a disk becomes totally corrupted (unfixable using a recovery system).

To create a restore dd tape:

```
# shutdown
# fsck -q /dev/dsk/6s0
# dd if=/usr/lib/uxbootlf.700 of=/dev/rmt/0mn bs=2k
# dd if=/dev/rdsk/6s0 of=/dev/rmt/0m bs=64k
```

To restore:

```
Select from menu:  b p2 ipl

ISL> hpux restore disk(scsi.6;0)
```

50722 11-10 116

Student Notes

In the event that a disk becomes totally corrupted or must be completely replaced, the recovery system cannot help you. The Series 700 **Restore Tape** is a different mechanism for recovering a corrupted root file system. Unlike the recovery system we just discussed, the restore tape does not fix the root disk, it *completely overwrites it* with a copy of the disk you made in the past. Anything on the root disk will be destroyed and replaced with what is on the restore tape.

The restore tape is created on a DDS/DAT drive.

The dd commands shown on the slide are used to build the restore tape. The first dd command is copying a boot area onto the tape. The second dd command is copying the entire root disk **image** onto the tape following the boot area.

Once you have your system installed and configured the way you like it, you can execute these two commands to build a new disk image. If your root file system should be destroyed beyond repair, you can boot the restore tape and copy the disk image from the tape back onto the disk. The disk will look exactly as it did when the tape was made using the dd commands. You may then have to recover a backup to completely restore your system.

To use the restore tape you must boot your Series 700 in attended mode by pressing Esc *once* when prompted after powering on the system. You will eventually see a menu like this displayed:

```
Device Selection    Device Path              Device Type
----------------------------------------------------------------------

P0                  scsi.6.0                 QUANTUM PD425S
P1                  scsi.5.0                 QUANTUM PD210S
P2                  scsi.3.0                 HP      HP35450A    -A

b)    Boot from specified device
s)    Search for bootable devices
a)    Enter Boot Administration mode
x)    Exit and continue boot sequence
?)    Help

Select from menu:
```

At this prompt you should enter: b p2 ipl Where p2 must be the identifier for your DAT drive.

At the ISL> prompt you enter the command to copy the disk image from the restore tape onto the specified disk. This command will restore a dd disk image to a scsi disk at address 6:

```
ISL> hpux restore disk(scsi.6;0)
```

Make sure the specified disk is the same size as the disk the recovery system was made from. Since we are working with a disk image the sizes must match.

Once the copy completes, you will be able to boot from the specified disk and you will see the disk look exactly as it looked when the image was made. If you had made changes to your system since the restore tape was made, you can simply recover the necessary changes from your most recent backups.

11-10. SLIDE: Series 700 Restore Tape — Instructor Notes

Purpose

This method is unique to the series 700 and provides a quick way to make a full physical backup of the root disk.

Key Points

- This is not a recommended way to perform full system backups. You cannot recover just one file from the restore tape. It is a disk image and you can only recover the image.
- This is much easier to do than building and using a recovery system but it does not have the flexibility of the recovery system.

11-11. WORKSESSION: Review Questions

1. What is the difference between a Series 700 recovery system and restore tape?

2. What is the purpose of the `/usr/lib/cron/cron.allow` file?

3. Can cron be used to bring the system to a single user run-level, perform a system backup then reboot back into the default run-level?

4. Why would you want to remove a software subsystem from your kernel?

11-11. WORKSESSION: Review Questions

Instructor Notes

1. What is the difference between a Series 700 recovery system and restore tape?

Answer:

The restore tape will completely rewrite the root disk whereas the recovery system will give you tools so you can attempt to fix it without destroying the current contents.

2. What is the purpose of the `/usr/lib/cron/cron.allow` file?

Answer:

This file contains the list of users that are allowed to use `crontab`.

3. Can cron be used to bring the system to a single user run-level, perform a system backup then reboot back into the default run-level?

Answer:

NO. Once the system is changed to a single-user run-level all processes including your job and the cron daemon are gone. The system would sit there in single-user mode until someone manually brought it back up.

4. Why would you want to remove a software subsystem from your kernel?

Answer:

If you loaded a software package like the X.25 communications software on your system, the kernel would be configured to support X.25. If you remove the X.25 software using `rmfn` you should also take it out of the kernel to reduce its size.

11-12. LAB: Lab Exercises

Directions

You can all work simultaneously on the `cron` exercises if you login using different names. The exercises for kernel reconfiguration and the recovery systems must be done as a group.

1. Schedule a job to run every 15 minutes that displays the date to your screen.

2. Run SAM and enter the **`Kernel Configuration ->`** menu. Go to the `View/Modify I/O Configuration ...` screen and note the name of the driver used to support SCSI DAT tapes.

3. Press **`Exit Task`** to return to the "*Kernel Configuration*" screen then select **`Modify Operating System Parameters ->`**. Try to find the **`maxusers`** and **`maxdsiz`** parameter. Use the **`Help`** facility to determine what they are used for. Look around at some of the other kernel tunable parameters and see if you can tell what they are used for.

4. Series 300/400/700 Only: If the hardware is available in your classroom, make a recovery system using the `mkrs` command. If time permits, boot the recovery system and try a simple task like removing the root password.

11-12. LAB: Lab Exercises

Instructor Notes

1. Schedule a job to run every 15 minutes that displays the date to your screen.

Answer:

The cronfile would look something like this:

```
0,15,30,45 * * * * date > /dev/your_tty_name
```

You can find *your_tty_name* using the `who` or `tty` commands.

You then submit your cronfile to cron using `crontab`:

```
$ crontab cronfile_name
```

2. Run SAM and enter the `Kernel Configuration ->` menu. Go to the `View/Modify I/O Configuration ...` screen and note the name of the driver used to support SCSI DAT tapes.

Answer:

The driver name is `scsitape` on Series 300/400/700 and `tape2` on the Series 800.

3. Press `Exit Task` to return to the "*Kernel Configuration*" screen then select `Modify Operating System Parameters ->`. Try to find the `maxusers` and `maxdsiz` parameter. Use the `Help` facility to determine what they are used for. Look around at some of the other kernel tunable parameters and see if you can tell what they are used for.

Answer:

`maxusers` is on the "*Miscellaneous Parameters*" screen. It is used to scale many other system parameters. It has nothing to do with how many users can log into the system. `maxusers` is one of the most commonly changed kernel parameters.

`maxdsiz` is the maximum amount of data space a process can allocate. Applications that use large data sets (like 3D CAD or large database applications) often need a larger value for `maxdsiz`. This parameter is found on the "*Process Parameters*" screen.

4. Series 300/400/700 Only: If the hardware is available in your classroom, make a recovery system using the `mkrs` command. If time permits, boot the recovery system and try a simple task like removing the root password.

Answer:

See the "Series 300/400/700 Recovery System" section of this module for the details.

Appendix A — Installing HP-UX

About This Appendix

The process of installing an HP-UX system is extremely simple. You are prompted throughout the process to customize your installation. There are, however, a few areas that can be expanded on.

This appendix describes how to get your HP-UX system installed for the first time (or reinstalled if the root file system was somehow destroyed).

The exact installation process will vary from one machine type to another and from one release of the operating system to another.

Reference

The *Installing and Updating HP-UX* books for each type of machine will be your primary reference during a system installation. You should read this book carefully before starting an installation.

Overview of Module A

This audience should not be doing their own system installations. There are times, however, when it may necessary to do so.

The material in this appendix is detailed and can tend to be confusing. A demonstration of the installation procedure usually goes over much better than this appendix.

If you decide to do a demonstration, you will probably still want to cover things like `rmfn` at the end of this appendix.

Disclaimer: The material presented here and in the update appendix is too detailed and often too specific to the workstation audience. Cover this only to satisfy the basic needs of the class, try to do a demo instead of covering this appendix and refer the students to the *Installing and Updating HP-UX* books.

A-1. SLIDE: HP-UX Installation

HP-UX Installation

- Destroys any data that previously existed on the root disk.
- Constructs an HP-UX file system on the root disk.
- Copies subsystem files from the installation media to the root disk.
- Builds the initial HP-UX kernel on the root disk.
- Creates a login for the system administrator.

HP-UX installation devices include:

Series 300/400:	Cartridge tape and CD-ROM
Series 700:	DDS format DAT and CD-ROM
Series 800:	DDS format DAT and CD-ROM

50722 A-1 117

Student Notes

Once you have properly installed your hardware, you are ready to install the software. Whenever a system installation needs to be done, the HP-UX system administrator is the only individual who should be permitted to do it. A system installation is required when a new system is ordered without pre-loaded software or "instant ignition".

A system installation may need to be done if there is a catastrophic failure with the system. Installation assumes that there is no data of value existing on the disk.

Installation is not required when a new release of HP-UX is offered, usually an update is performed.

During our discussion of system startup it was mentioned that attended boot procedures allow you to boot from a device. Installation requires that you have the appropriate "install media" available to your system and that you perform an attended boot from it.

Once you have booted from the "install media" you must answer questions concerning the layout of your system disk. The install utility does a good job of holding your hand and prompting for answers.

A-1. SLIDE: HP-UX Installation — Instructor Notes

Purpose

To de-mystify the system installation process and prepare the students to interact with the install utility.

Key Points

Point out that once you have booted the system from the install media the purpose of the utility is to prepare the destination root file system by initializing it, creating the boot area, reserving primary swap space and copying a minimal HP-UX file system to it.

Teaching Points

Don't get lost here on the differences of the architectures or the various media instead concentrate on the similarity of the task at hand.

Transition

Before you start an install you need to prepare for the questions asked.

A-2. SLIDE: Before You Start an Installation

Before You Start an Installation

- Verify that system hardware has been correctly installed.
- Read the *README BEFORE INSTALLING HP-UX* notice.
- Determine:
  - The hardware address of the installation device and the root disk.
  - The amount of swap space needed.
  - Any required codewords and hardware IDs.
  - Which filesets and partitions you wish to install.
  - Customizing information (time zone, time, host name, IP address).

50722 A-2 118

Student Notes

Before you start an install, do the following:

- Verify that system hardware has been correctly installed.
- Read the *README BEFORE INSTALLING HP-UX* notice. If there was an installation manual supplied with your installation media, read it as well. Other excellent manuals for installation are *Installing and Updating HP-UX HP 9000 Series 300/400* or *Installing and Updating HP-UX HP 9000 Series 700.* or *Installing and Updating HP-UX HP 9000 Series 800.*
- Determine the hardware address of the installation device and the root disk. Determining these addresses is not absolutely necessary as the boot and installation processes will search for possible sources and destinations. However, knowing these addresses will help you confirm the correct device addresses as the installation program runs.

- Determine the amount of swap space needed. This information should be determined from the installation notes that came with your application software. The *README* and *Installing and Updating* books can also give you a rough idea of how much swap space is necessary.
- Determine any required codewords and hardware IDs. Starting at the HP-UX 8.0 release, with the introduction of CD-ROM software distribution, codewords have been added as a security measure to the installation process. The codewords are based on specific software products and are tied to specific computer system hardware IDs. Codewords will be supplied on a software product certificate that you receive with your distribution media. The hardware IDs are derived from the specific computer system that will run the software you are installing. At this time, only installations using CD-ROM media require the use of codewords. Codewords and hardware IDs are covered in detail in the update appendix.
- Determine which filesets and partitions you wish to install. If you are not planning to install everything from the software bundle that you ordered, determine which filesets and partitions you wish to install. A partition is a group of filesets, and a fileset is a group of files. The referenced documents will give you complete lists of the partitions and why you might want to load each one.
- Determine customizing information (time zone, time, hostname, IP address). If you plan to use networking, you will need to know a unique hostname for your system and a unique Internet Protocol (IP) address for your LAN interface. This information may not be available at the time you install the system, so there are defaults provided for you. Hopefully, time zone and time should be obvious and easily obtainable at install time.

A-2. SLIDE: Before You Start an Installation — Instructor Notes

Purpose

To inform the students of preparatory steps they should take before beginning their installation.

We also introduce concepts of software partitions, filesets, and files early in this appendix. Frequent reference is made to this concept throughout the appendix to these terms.

Key Points

Before starting your installation, write down the key information that will be required during the process:

- Hardware address of your installation device and the root disk which you plan to put the operating system on.
- Codewords and hardware IDs if using CD-ROM installation
- Swap space required by HP-UX and your applications.
- If planning to use networking, a host name for your system and an IP address. This information may not be available.
- Determine which filesets and partitions to load.
- Determine the time zone, current time, host name, IP address.

Hardware Addresses

Refresh the students' memories with

Series 300/400:	Disks are identified by select code and bus address.
Series 700:	Disks are identified by the card slot number and the bus address.
Series 800:	See the appendixes of the *Installing and Updating* manual for details on the different types of hardware addresses.

Amount of Swap Space is Required.

Determining the amount of swap space required is not as easy and straight forward as one would hope. You should see the HP-UX System Administration Tasks Manual for more details. You must read the manuals and `README FIRST` notices supplied with each of your applications software and subsystems to determine the necessary swap space.

Some rules of thumb for swap space:

- For an **HP-UX Runtime System** (this is the new name for the HP-UX Application Execution Environment - AXE): Use default swap values for your root disk if a few small applications are to be run. Larger and swap intensive applications will require more swap.

- For a typical HP-UX workstation: Use default swap values for your root disk if you plan to use the X11 Window System, and subsystems such as SoftPC, networking, and BASIC/UX. If large and swap intensive applications are to be run (for example, LISP or ME10) you will require more swap space.
- For an **HP-UX Clustered Environment** (This is the new name for the HP-UX Diskless Cluster): Approximately 300 megabytes of file system space will be required for a cluster server. The swap space should include 20 megabytes for each `cnode` (client system) being served. This amount may need to be larger depending upon the application being run. Remember as swap increases, file system size decreases for a fixed size disk.

Remember that primary swap space, once reserved on a disk, cannot be reclaimed without rebuilding the file system. This may require re-installing HP-UX on your system disk or in the case of auxiliary swap disks simply rebuilding the file system. In both cases restoring files from a backup will be necessary. There is a lot of time and effort required in either case.

Fortunately at the HP-UX 8.0 release it is very easy for you to add swap space after the system is installed. There is no need to rebuild the kernel in most cases when adding disks to provide swap space. Adding `dynamic` swap space (a new feature at 7.0) has become more flexible at the 8.0 release.

Teaching Tips

As a relevant example for unwanted software, you may have included a language translator (for example, Pascal) in your purchase (for instance in a software bundle), but do not plan to use it on the system due to a change in plans.

A-3. SLIDE: Steps to Install the HP-UX Operating System

Steps to Install the HP-UX Operating System

1. Turn on power to the disk drive, installation device, monitor, and expander (if applicable).
2. Insert the media labeled *INSTALL* into the installation device. If using tape, wait until the "busy" light goes off.
3. Turn on SPU.
4. Boot in attended mode.

50722 A-3 119

Student Notes

Using CD-ROM Install Media

Before inserting an *INSTALL* CD-ROM disk into the CD-ROM drive, make sure that the drive is powered ON and the front panel lights are out.

You should not have to force the disk carrier into the drive, but you will need a firm, steady push. Hold the disk carrier at the end labeled "Compact Disc", and push the carrier evenly and firmly into the drive until it bottoms-out. Then you should push it a bit further, release it and the carrier will pop back out partially, leaving the CD-ROM inside the drive. Pull the now empty carrier fully out, and set it aside, you will need it later to extract the CD-ROM in the drive.

The new CD-ROM drives use a carrier with a flip-open lid. This type of carrier simply inserts into the slot and the drive pulls the carrier and its CD in together.

Using Cartridge Tape Install Media

Before inserting an "install" tape into the cartridge tape drive, check to make sure that it is write-protected or "safe". The arrow in the top left corner should be turned to point to "SAFE."

Using DDS Format DAT Install Media

Before inserting an "install" tape into the tape drive, check to make sure that it is write-protected or "safe". The white tab should be at the open position.

A-3. SLIDE: Steps to Install the HP-UX Operating System

Instructor Notes

Purpose

To detail the initial steps necessary to install HP-UX. You may think that some of this is very basic, but if the students have never installed a system before, or don't have much experience with hardware, they will have no idea about the order in which they should power up their components. If they have hardware experience, this will be a good review. So, cover all the steps, and encourage experienced users to help others in the labs that require this knowledge.

Turn on the disk drive, tape drive, I/O expander, and monitor in the order described on the slide. There is a "method" here, you should follow it.

Some disk models may have a different color indicator light and a different manner of indicating "busy/access", so consult the reference manual for your drive. The references in these notes to ON and OFF are based on typical CD-ROM, tape and disk devices, yours may differ.

Remind students to check that the install cartridge tape or the DDS format DAT is write protected before inserting it into the tape drive. There is no equivalent for CD-ROM disks since they are inherently read only.

After inserting the cartridge tape, wait until the busy light goes off and **stays off** on the tape drive before continuing. The indicator light will blink on and off for several minutes.

Transition

After the "install" utility has done its job you have a minimal system on the root disk and now we need to update to a fully loaded system

A-4. TEXT PAGE: Rebooting From System Disk

Screen Facsimile

```
Copyright 1989
CONSOLE is 98644 at sc 9,
1 port(s)
 . . .
Internal HP-IB Interface - system controller at select code 7
  Parallel poll interrupts enabled.
HP98644 RS-232C Serial Interface at select code 9
HP98265A (SCSI Interface) sync 2.67 MBytes/sec; parity
enabled at select code 14
HP98643 LAN/300 Link at select code 17 ignored; interface driver not present
real mem  = 4182016
 . . .
Initializing . . .

        Ensure that the install media unit has been removed and an
        update media unit is online and prepared for reading.

                --- Press "Return" to continue ---
```

Student Notes

As the system reboots from the root disk, you will see now familiar messages on the screen. Next you will be prompted to remove the INSTALL medium and insert the first UPDATE medium into the installation device.

You must first unload the INSTALL medium. Next you should load the first UPDATE medium. Wait for the busy light to go off on the installation device. If this is a cartridge tape drive, the busy light should go out and remain OFF (it will blink on and off for several minutes). With a DDS format DAT, the tape will take about 25 seconds to load and the upper light will flash green. The upper light will be yellow for a write-protected tape. This is the case during an installation. (For a write- enabled tape, the upper light will be green.) With a CD-ROM the drive will be ready in about five seconds, the ready light will blink once very quickly and remain ON. When the installation device is ready, press (Return). You will see the menu shown on the next slide.

Warning **You MUST load the UPDATE media unit at this point and wait for it to become ready. You must first unload the INSTALL medium to do this. If you fail to swap the media as directed at this point when using the CD-ROM installation media, or continue by pressing (Return) before the CD-ROM drive is ready, you cause an unrecoverable error. You can continue for a few steps but will reach a point where**

you are unable to complete the installation process. If you are using tape, the program will allow you to recover from this error.

A-4. TEXT PAGE: Rebooting From System Disk — Instructor Notes

Purpose

To explain what happens in the installation process following the automatic reboot. This is the start of the "Update" phase of the installation.

Key Points

- The automatic reboot uses the root disk instead of the installation media unit we were using initially.
- The operating system that starts running is from the HP-UX core filesets just loaded. The core file sets contain the kernel, a shell, system libraries and a few essential system administration tools. This was covered in the execution trace printed in the Student Notes.
- You are prompted to replace the install media with update media.
- You MUST load the **UPDATE** media unit at this point. Otherwise you will cause an unrecoverable error when using CD-ROM install media.
- You must wait for the tape unit's busy light to go off (and stay off) and then press [Return]. If a CD-ROM, wait about ten seconds after the green light blinks once.

Transition

You have now loaded the first UPDATE medium and waited for the installation device to become ready. When you pressed [Return], you have entered into the "Update" portion of the installation process, you will see the menu shown on the next slide.

A-5. SLIDE: Update Main Menu

Update Main Menu

```
INSTALL                        Main Menu

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

     Source: Tape Device            Destination: Local System
             /UPDATE_CDROM                       /
     -----------------------------------------------------------

          Select All Filesets on the Source Media
          Select Filesets for a Minimum System ->
          View/Select Partitions and Filesets ->

          Enter Codeword ->

          How to Use Install

 Help         Shell   Select                               Exit
                      Item                                 Install
```

50722 A-5 120 ©1992 Hewlett-Packard Co.

Student Notes

The next phase in the installation process is to add more filesets to the existing HP-UX file system using a running version of HP-UX. This menu is actually the first menu of the update process. Details on update are covered in the update appendix.

A-5. SLIDE: Update Main Menu

Instructor Notes

Purpose

Explain where update starts.

Point out that update is covered in the update appendix.

Don't cover details on:

```
Select Filesets for a Minimum System ->
View/Select Partitions and Filesets ->
```

Transition

After the update portion is complete, install customizes the system.

A-6. SLIDE: Customizing the System

Customizing the System

You will need to provide the following:

- Continent where your system will be located
- Time zone
- Current date (day, month, and year)
- Current time of day (hours and minutes) using 24 hour clock
- The unique network host name assigned to your system
- The IP (Internet Protocol) address assigned to your system

50722 A-6 121

Student Notes

The rebooting at the end of file loading is nearly the end of the INSTALL process. You must interact with a special program, `/etc/set_parms` that runs at the first reboot. Completing that interaction represents the end of the installation process, but not the end of your responsibilities for the installation.

Eventually as the system reboots, you will see a `Console login:` prompt. Before the prompt appears you will be required to answer a series of questions to set up system customization files. Once you get the console login prompt, you must log in as `root` (no password is required) and perform additional postinstallation procedures (covered in the next appendix).

Following the reboot, you will see system startup messages on the console as before. The messages will be interrupted just after the `/etc/rc` script starts, with a series of questions from `/etc/set_parms` that you must answer to customize your system for your locale and time of day.

Before you respond you will need to know the following data:

Continent where your system will be located.
Time zone
Current date (day, month, and year).
Current time of day (hours and minutes) using 24 hour clock
The unique network host name assigned to your system.
The IP (Internet Protocol) address assigned to your system

You can use the defaults for the network host name and the IP address if they are not known.

The messages you will see are listed below. Your first question will be:

```
/etc/rc:

############################################################
The following procedure will allow you to set the time zone.

Select your continent from the following list:
              1. North and Central America
              2. South America
              3. Europe
              4. Africa
              5. Asia
              6. Australia, New Zealand

Enter the number corresponding to your continent (1-6) -> 1
```

The correct answer for the USA is 1. Next you will be asked:

```
#############################################################
Select your time zone from the following list:

1. Newfoundland Std. Time (NST3:30NDT) |  7. Mountain Standard Time (Arizona)
   Newfoundland Daylight Time          |     (No daylight savings time.)
                                       |
2. Atlantic Standard Time (AST4ADT)    |  8. Pacific Standard Time (PST8PDT)
   Atlantic Daylight Time              |     Pacific Daylight Time
                                       |
3. Eastern Standard Time (EST5EDT)     |  9. Yukon Standard Time (YST9YDT)
   Eastern Daylight Time               |     Yukon Daylight Time
                                       |
4. Eastern Standard Time (US:Indiana)  | 10. Aleutian Standard Time (AST10ADT)
   Central Daylight Time (EST5CDT)     |     Aleutian Daylight Time
                                       |
5. Central Standard Time (CST6CDT)     | 11. Hawaii
   Central Daylight Time               |
   Central American countries          | 12. Unlisted time zone
                                       |
6. Mountain Standard Time (MST7MDT)    | 13. Previous menu
```

```
    Mountain Daylight Time                    |

  Enter the number corresponding to your time zone (1-13) ->5
```

The correct answer will obviously vary depending upon your site. Next you will be asked a series of questions about the date and time:

```
  #####################################################################

  You will be prompted for the date and time.  Please enter all values
  numerically, for example January is 1.  The values in the parenthesis
  give the acceptable range of responses.

  Please enter the month (1-12) -> 1
  Please enter the day of the month (1-31) -> 25
  Please enter the hour (using 24 hour time) (0-23) ->15
  Please enter the minute (0-59) -> 16
  Please enter the last two digits of the year (70-99)->91

  The date and time have been set to: Fri Jan 25 15:16:00 CST 1991
```

Your response of course varies, but should be relatively accurate. Next the system asks you for networking information. You should have already selected a host name and an IP Address for your system. You can use (Return) as a default.

```
  For the system to operate  correctly  you must assign it a
  system name.  It should be 8 characters or less and may
  contain any alphabetic character, 0 through 9, -
  or _ (hyphen/dash and underscore).

  The first character must be alphabetic. Uppercase alphabetic
  characters are allowed but not recommended.  Entering nothing
  followed by return will enter the default name of "unknown".

  Enter the system name  -> training
```

If you provided an acceptable host name, you will then be prompted for a unique IP address. As you can see reasonable defaults are provided:

```
################################################################
If you wish networking to operate correctly, you must assign
the system an Internet Protocol (IP) address.  This  address
must be assigned  by your local  network  administrator.  If
you do know your IP address you may press  return to use the
default value (127.0.0.1).  An IP  address  should  consist
of four numbers separated  by periods.  The numbers  should
be between 0 and 255 inclusive.  For example: 255.32.0.10

Enter your Internet Protocol address -> 192.6.241.71
```

At this point You must simply wait for the system to complete the rebooting process. Your system will become very busy for a few minutes then complete the boot process.

At the end of the process your screen will have a number of

messages such as:

```
Starting up standalone system
Is the date Fri Jan 25 15:16:35 CST 1991 correct? (y or n, default: y) y
```

You can reply or let it time out and take the default of "y". This will be followed by more boot up messages. Many will indicate some form of error. They can be safely ignored at this time. You will see messages similar to the following:

```
/etc/nettl: !: not found
ifconfig: socket: Protocol not supported
lanconfig: socket: Protocol not supported
ifconfig: socket: Protocol not supported
/bin/nodename: system feature not installed
Network Link started
System message logger started
    starting NFS networking
        NIS domain name not set
        /etc/portmap
    Network Information Service not started.
ARPA/Berkeley daemons started: inetd
Network NS Services started
Network management daemons started: socket: Protocol not
supported

NETWORKING started.
cron started
starting the ptydaemon
starting the vtdaemon
NOTE:  Files in /tmp:
```

```
total 236
crw-r--r--   1 root     other      4 0x070000 Nov  7 17:00 _chk_for_rombo
-rw-r--r--   1 root     other        765 Jan 25 15:19 nettlgen.log
-r-xr-xr-x   1 bin      bin         9096 Nov  6 14:54 rgb.800
-rw-r--r--   1 root     other      32480 Nov  7 23:24 rmfn.log
-rw-r--r--   1 root     other         24 Nov  7 23:26 update.cleanup
-rw-r--r--   1 root     other      75858 Nov  7 23:26 update.log
/etc/auditrc: This file must be edited to activate audit subsystem
Fri Jan 25 15:19:40 CST 1991

Console Login:
```

Finally at the console login, you can become the superuser by typing "root" (Return), and proceed to the post installation procedures as described in the next appendix.

A-6. SLIDE: Customizing the System — Instructor Notes

Purpose

To show the screen that comes up while the installation is actually loading partitions and filesets. Also to show the final steps of system customization which are now part of the Install process before the first "Console login:" prompt appears..

Key Points

The process changes in an important way for installations based on codeword protected media (that is, CD-ROMs) at 8.0 Take the time to be sure the students understand that the process is similar between the two, differing only in codeword management.

After all the filesets are loaded, your install is complete. The system will reboot. Log in as `root` and perform your postinstallation procedures.

A-7. SLIDE: Post Installation Procedures

Post Installation Procedures

- Execute basic procedures to ensure system integrity.
- Customize the system.
- Create a recovery system on a Series 300/400/700.

50722 A-7 122

Student Notes

Once installation is finished, the job of system administration is just beginning a new facet. Before going into the daily care of an HP-UX system there are some steps that need to be taken:

- Execute basic procedures to ensure system integrity. To guard against sabotage and recover from failures, there are some safeguards that can be employed.
- Customize the system. The system needs to be customized for such things as its name and time zone. Files executed at login will need modification. Unwanted files will need to be removed, and man pages will need to be formatted.
- Create a recovery system on a Series 300/400/700. A recovery system needs to be created on a Series 300/400/700. A recovery system is used in case the system is unbootable due to a corrupt root disk. A recovery system provides tools to repair a corrupt root disk. On a Series 600/800 the support tape is used in such situations.

A-7. SLIDE: Post Installation Procedures

Instructor Notes

Purpose

Introduce the types of procedures that need to be performed after installation of the software.

A-8. SLIDE: Procedures to Ensure System Integrity

Procedures to Ensure System Integrity

- Unload all *UPDATE* tapes and store in a safe place.
- Assign a root password.
- Check install/update information in `/tmp` and "/etc/newconfig".
- Print superblock information from `/etc/super_blocks` or "/etc/sbtabs".

50722 A-8 123

Student Notes

Unload and Store all UPDATE Tapes

As soon as possible following the installation, unload all the UPDATE media used in the installation, and along with INSTALL media, store them in a safe place. Failure to do this may result in a security problem. A person in possession of these tapes who has physical access to your system can destroy data and compromise your system security.

Assign a Root Password

To ensure that your system is secure, assign a root password immediately. Use the `passwd` command to set a superuser password, as shown below. The passwords you type will not be echoed to the terminal, and you must enter them twice.

```
# passwd root
Changing password for root
New password: ********
Re-enter new password: ********
#
```

You should not write this password down anywhere, as it would be more likely to be discovered by someone else. However, should you forget this password, there is no easy way to change it. On a Series 600/700/800 system, you must boot in single user mode to change the password. (See separate appendixes on boot up for details.) On a Series 300/400 and also a Series 700, the recovery system can be booted and then the password can be changed. (See material at the end of this appendix for details.)

Check install/update information

After the installation you must check the install/update information logged during the process. To do this you should check the log files in the `/tmp` directory and check files placed in the `/etc/newconfig` directory.

The log file `/tmp/update.log` records a log of any errors or warnings that happened during the install. A careful reading of that file is important to determine if the installation was correctly carried out. The entries entered in this file were discussed in the appendix on "Installing HP-UX", which you can consult for the meaning associated with errors, warnings, and "Notes" posted to the log.

The `/etc/newconfig` directory contains some files that hold information about your new release. The file `/etc/newconfig/`*UpdateInfo* gives you information on the release you are installing. A careful reading of this file is important prior to performing an update, and less critical to an installed system, unless you need to know what changed since the previous release of HP-UX.

The `/etc/newconfig` directory contains files that were referenced in messages from the `update.log` file mentioned above. These are usually copies of files that you will need to copy into another directory, based on `update.log` "Note" messages.

Print Superblock Information

It is important that you know where the backup copies of your primary super block are located. On a Series 600/800 the file `/etc/super_blocks` contains this information. On a Series 300/400/700, the file `/etc/sbtab` contains this information. Print a copy of the appropriate file for your system, and save the printout where it can be retrieved. A good place for keeping this important data is in the binder you use for installation and maintenance history and for a system log.

A-8. SLIDE: Procedures to Ensure System Integrity

Instructor Notes

Purpose

Go over some simple procedures to ensure system integrity.

Key Points

See list on graphic.

A-9. SLIDE: Checking Up on update

Checking Up on update

- Verify time is correct by examining the contents of /etc/src.sh and /etc/src.csh.
- Remove files not removed during update.

50722 A-9 124

Student Notes

Verify Time is Correct

Several startup customization files are "sourced" to customize the system configuration at boot time (set the time zone and host name). You may choose to verify their contents at this time; however, at HP-UX 8.0 these are edited for you as part of the menu that comes up after the installation utility has installed the software.

To verify customized data in startup files, check two files created at the end of the installation process, /etc/src.sh and /etc/src.csh. As of the HP-UX 8.0 release, you no longer have to edit certain system files to ensure that the time zone environment variable (TZ) and the system hostname are correctly set. TZ is referenced in:

- /etc/rc

- /etc/profile
- /etc/csh.login
- /etc/powerfail on the Series 600/800

The two files /etc/src.sh and /etc/src.csh were created in the installation process to set the TZ variable and hostname for you at startup time. They are now "sourced" by the /etc/rc script at boot time and the system-wide profiles (/etc/profile and /etc/csh.login) as users login.

The contents of /etc/src.sh should be something like:

```
TZ=CST6CDT; export TZ
SYSTEM_NAME=training ; export SYSTEM_NAME
```

The contents of /etc/src.csh should be something like:

```
setenv TZ CST6CDT
```

The time zone and host name values will obviously vary depending upon your location. The values here are typical for the central USA timezone. To recreate these two files you can re-run the interactive program /etc/set_parms.

The files /etc/src.sh and /etc/src.csh are "sourced" inside shell scripts at startup time to set the environment variables used in /etc/rc for setting the network name of the system and the time zone for formatting the date command output.

The program /etc/set_parms will not reset the system clock for you if the file /etc/installtime exits. You should not delete the file /etc/installtime as its time stamp is a reference indicating when your system was installed.

For more information on time zones see the file /usr/lib/tztab.

Remove files not removed during update

During the update portion of the installation, the install or update process may be unable to remove files from the file system. This occurs when there is an attempt to remove a file from the disk when a process is running that corresponds to the file. As an example, it may well be that the file to be removed was /etc/update, and it was also executing. The program would rename it to /etc/#update, and unload a new /etc/update from the update media.

The update process will usually rename that file with a name formed by placing an "#" in front of the old filename. You should remove all such files after the installation or update completes. The list of such files is in /tmp/update.cleanup and can be removed with the command:

```
# rm -rf `cat /tmp/update.cleanup`
```

A-9. SLIDE: Checking Up on `update` — Instructor Notes

Transition

Other post installation steps include removing unwanted optional filesets, reformatting the manual pages, setting the system clock, and making a recovery system.

Next we look at removing unwanted optional filesets. This is done with the `/etc/rmfn(1m)` command.

A-10. SLIDE: `rmfn` — Removing Optional Software

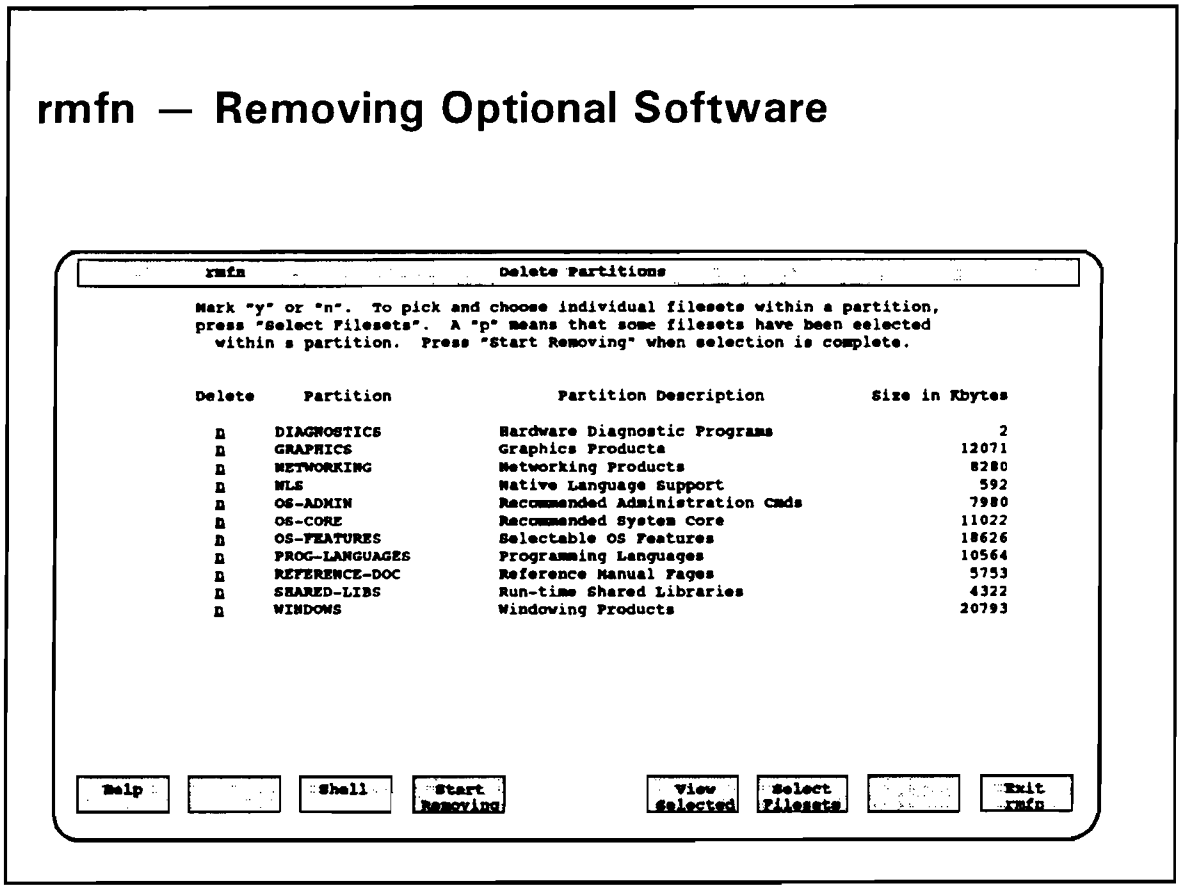

50722 A-10 125

Student Notes

It is recommended that when you do your installation, you install all products on your distribution media. After everything is installed, you can remove the filesets that you do not want. The filesets that are installed on your system are listed in the directory `/etc/filesets`.

Starting at the HP-UX 8.0 release, the `rmfm` command can be used to remove unwanted optional filesets. The `sysrm` command was used for this purpose previously, and is now obsolete. The `sysrm` command had a serious limitation: it allowed you to remove a fileset needed by another product's files. The interactive utility `rmfn` has a much better understanding of fileset dependencies, and will prevent this from happening accidentally.

To use the interactive utility `rmfn(1m)` command to remove the unwanted optional filesets invoke it with:

```
# rmfn
```

You will have to wait a short time while the filesets that are loaded on your system disks are scanned for size and dependency information. During the delay you will see the following message on the screen:

```
Getting fileset information ...
```

Then the screen shown in the slide will appear, listing the partitions loaded on your system, along with their sizes. At this point you can select a partition for deletion by changing the n to a y in the "Delete" column. When you select a partition for removal it will again take some time for the system to traverse the list of dependencies for the filesets in the partition you select for removal. This is usually ten or twenty seconds of intense disk activity before the y is posted to the screen.

Note

Do not press y twice, thinking that it did not recognize the first one, because this may cause you to select another partition you did not intend. The program moves the cursor to the next line after a selection is made. A second affirmative response will select the next partition.

After you have selected a partition, the softkey `View Selected` can be used to determine exactly what filesets will be removed.

A-10. SLIDE: rmfn — Removing Optional Software

Instructor Notes

Transition

The screen that appears after you press the softkey `View Selected` lists the filesets to be removed. This is the screen on our next slide.

A-11. SLIDE: rmfn — View Partitions/Filesets Selected

rmfn — View Partitions/Filesets Selected

```
rmfn                    View Partitions/Filesets Selected

       These are the Partitions/Filesets that have been selected for removal.
       Selections can be changed at this time by entering "y" or "n".
       Press "Partition Screen" to return to the Delete Partition screen.

                                         Total in Kbytes:              592

  Delete        Partition               Fileset               Size in Kbytes

    y       NLS                     AMERICAN                             11
    y       NLS                     ARABIC                                7
    y       NLS                     ARABICW                               7
    y       NLS                     BULGARIAN                             6
    y       NLS                     CFRENCH                              11
    y       NLS                     CHINESES                              6
    y       NLS                     CHINESET                              6
    y       NLS                     CZECH                                 6
    v       vvvvvvvvvvvvvvvv        vvvvvvvvvvvvvvvv          vvvvvvvvvvvvv

 Help          Shell   Start                                     Partition
                       Removing                                  Screen
```

50722 A-11 126

Student Notes

This is the screen you will see when you select a partition and press the softkey `View Selected`. While in this screen you may decline to delete any of the selected filesets by changing the "Delete" field (shown on the slide as y) to n. The same caution stated above applies to multiple key presses in this screen.

If a dependency exists for a fileset selected for deletion, you will get a pop-up window that warns you about it and offers you a chance to also delete the dependent fileset. You must answer y or n depending upon your requirements. If you answer negatively, neither the dependent fileset or the corresponding fileset you had marked will not be removed.

Once you have selected all the filesets you want to delete, you must press `Start Removing` softkey. You will be asked to confirm your choice to start the removal process. A pop-up window will appear. No softkeys will be defined, but the empty softkey labels will be displayed. The only possible response is y or n at this point. To confirm that you want to go on, press"y".

Of course the number and sizes of files you choose will affect the number of kilobytes to be removed. This is you last chance to back out. You must answer y to start removing files. Once you have confirmed your choice to start the removal process, a new window will appear to provide a status report on the removal process.

A-11. SLIDE: rmfn — View Partitions/Filesets Selected

Instructor Notes

Transition

This screen will appear to give you real time updates on the progress of the fileset and partition removal.

A-12. SLIDE: `rmfn` — Removing Filesets

rmfn — Removing Filesets

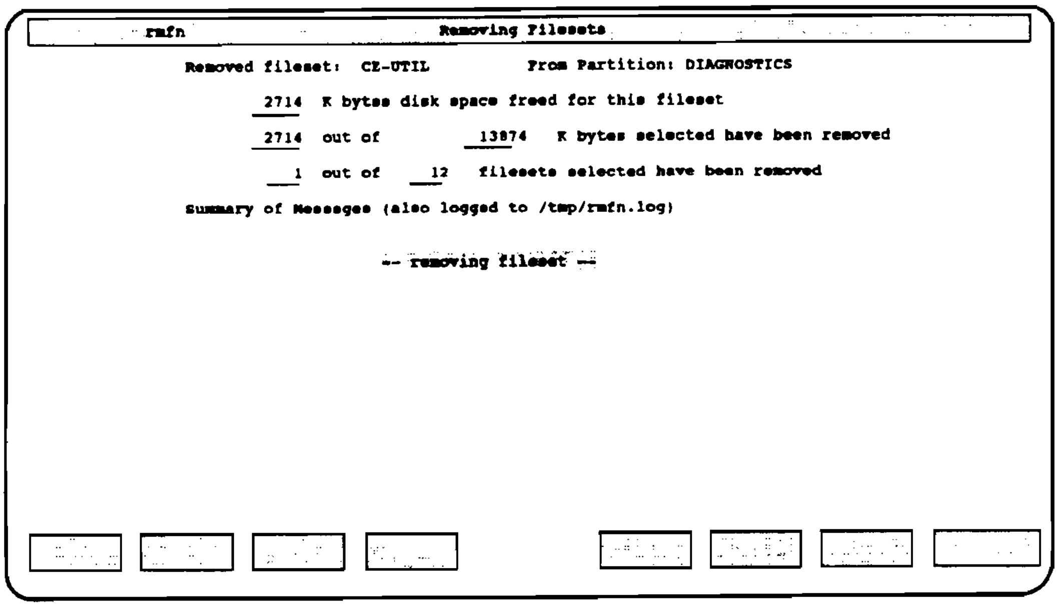

50722 A-12 127

Student Notes

This screen presents you with a continuous status report about the process of fileset and partition removal. It will display the progress and any error messages until the process completes. No softkeys are displayed while this screen appears.

The removal may take several minutes as individual filesets are removed. You will see a message flash at the bottom of the screen:

```
-- removing fileset --
```

At the end of the removal you will see a screen to indicate the end of the removal process. This screen will also report the disk space freed for this fileset. After this menu, select **`Exit rmfn`** if there are no other partitions and filesets to remove.

A-12. SLIDE: `rmfn` — Removing Filesets — Instructor Notes

Transition

Now that we have seen the process of fileset removal, we need to move on to other post installation steps.

Other post installation steps include reformatting the manual pages and setting the system clock and making a recovery system. Next we look at reformatting the manual pages.

A-13. SLIDE: Format Manual Pages

Format Manual Pages

- Create formatted versions of all manual pages and a complete `whatis` database

  `# /etc/catman`

- Create formatted versions of some sections of manual pages

  `# /etc/catman/1 1m 2`

- Do not create any formatted manual pages, but instead "build-as-you-use"

  `# man somecmd`

50722 A-13 128

Student Notes

Before we explain `catman` it helps if you understand how the manual entries are stored on the system. All the entries, both unformatted (source) and formatted, are stored in sub-directories under the `/usr/man` directory. Both unformatted and formatted entries can appear in compressed or uncompressed form. This gives us four possibilities.

	unformatted	**formatted**
compressed	/usr/man/man*X*.Z/*	/usr/man/cat*X*.Z/*
uncompressed	/usr/man/man*X*/*	/usr/man/cat*X*/*

X is the section: 1, 1m, 2, 3, 4, 5, 7, 9.
The unformatted entry is an `nroff` source.

catman (with no options) creates the formatted versions of the on-line manual pages from the compressed nroff source files. (Compressed files end in .Z.) When you invoke catman, each manual entry in the man*X*.Z directory is examined and those whose formatted versions are missing or out-of-date are created, compressed, and put it into the cat*X*.Z directory. Before invoking catman, it is best if you remove all of the catX directories (leaving only the catX.Z directories).

If you want to create uncompressed entries, make sure the catX directories exist and invoke catman -z. This command puts an uncompressed version of each entry in its appropriate cat*X* directory.

You can choose to format only certain sections of the manual pages by invoking catman *section*, where *section* is the number of the section you wish to format, for example sections 1, 1m and 2 as shown by:

```
# /etc/catman 1 1m 2
```

There are several user commands that have names used to reference the manual pages that appear in more than one section of the manual. Separate pages may exist for a user command (section 1), system administration command (section 1m), library routine (section 3), system call (section 2), file format (section 4), special file (section 7), or miscellaneous entry (section 5). Several examples are: mkdir(1), mkdir(2), mknod(1m), mknod(2), mknod(4), mount(1m), mount(2), and mount(3).

To read a specific manual page, type the man command with the section number as the first parameter. Several examples follow:

```
$ man 1 mkdir  OR $ man 2 mkdir
```

You can choose not to format any manual entries with catman. If you do not format entries, then the first time you try and access a manual entry you will get this message:

```
Reformatting entry.  Wait...
```

You will wait a short period of time before the manual entry appears on your screen. This is usually less than 10 seconds, but may be much longer for very large manual pages such as "sh". After the first time the page is read with man*some-command*, the entry for *some-command* is added to the appropriate cat*X* directory, and in subsequent accesses, the user does not have to wait. This is a "build-as-you-use" alternative. The system fills the cat*X* directories as the users access commands with man. For this to work, the cat*X* directories must exist. If they do not, you can create them with the following script:

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

Note

If you execute `catman` it requires several megabytes of disk space and several hours to complete. By doing so you get improved performance, but you sacrifice disk space. (To recover some disk space, you could delete the unformatted directories after using `catman`.)

Running `catman` allows you to use the `man -k keyword` command to locate documentation, executables and libraries by keyword lookup. A reference to the keyword in the one line description of the command is used as a search key for the "/usr/lib/whatis" database built by the `catman` command. All commands (files, man pages) with that keyword are also listed. This is like a "conceptual" cross reference to commands based on the keyword used.

A-13. SLIDE: Format Manual Pages — Instructor Notes

Purpose

Some administrators might want to create formatted entries of their manual pages. This slide explains how to do that.

Teaching Tips

Avoid the temptation of a discussion of `nroff` or `troff` as text formatters or simple desktop publishing tools. Keep the focus on post installation steps.

- Explain the directory layout of `/usr/man`.
- If users choose not to execute `catman` then the `cat`*X* directories must exist so that entries will be added to these directories as users use the `man` command. If these directories do not exist, the users can create them with the given script.
- Make sure to discuss the pros and cons of using `catman`. `catman` increases performance, but uses much disk space to store all the formatted versions of the manual entries.
- Mention that the multiple name problem may increase user requests for help when they get "funny" manual pages. An excellent compromise is to always format all the user commands in sections 1, if some users expect to read system calls (section 2) and library calls (section 3) manual pages.
- It usually takes 5 to 6 hours to run `catman`. Recommend that the students run the command in the background at night. You might type this just as you leave for the day:

```
# nohup nice /etc/catman &
```

- Running `catman` allows you to use the `man -k  keyword` command to locate documentation, executables and libraries by keyword lookup. A reference to the keyword in the one line description of the command is used as a search key for the "/usr/lib/whatis" database built by the `catman` command. All commands (files, man pages) with that keyword are also listed. This is like a "conceptual" cross reference to commands based on the keyword used.
- There is no relationship between catman and the `whereis` command. The `whereis` command is useful to see what kinds of man pages exist for a given name, but it does not require that `catman` have been run. It is a Berkeley derived command and may not be familiar to HP-UX users who came from a System V environment.
- You might want to mention the `fixman` command. This shell script processes all ordinary files under `/usr/man/cat`*X* to remove all character and backspace pairs. Such pairs usually exist to cause overstriking or underscoring for printer output. They slow down `man`, and use up more disk space. `fixman` should be run after `catman` to rebuild all cat-able manual entries.

Lab possibilities

In most class room setups the students will not be able to do an actual installation, on the Series 600/800 there is a simulation program that gives the student some idea of how an install goes. If you have an extra system available it is nice to do a class demo of an install.

Appendix B — Updating HP-UX

About This Appendix

This appendix covers updating your system software. Typically, you will update your system when Hewlett-Packard releases a new version of the operating system or when new application software versions are released. You can also use the update process to add new subsystems (or other applications) to your current system.

Note

This appendix applies to all HP 9000 systems running only the HP-UX 8.X release. This information is very likely to change when updating to HP-UX 9.0. To be absolutely certain that you follow the correct procedures for updating HP-UX find the *Installing and Updating HP-UX* book for your type of system and the newest software release.

Overview of Module B

See the instructor notes at the beginning of Appendix A.

Stress that it is very likely that the students are already running HP-UX 8.0 and this material may not directly apply to HP-UX 9.0. The general procedures will be the same but the specific commands and methods could vary greatly.

B-1. SLIDE: Installing Versus Updating HP-UX

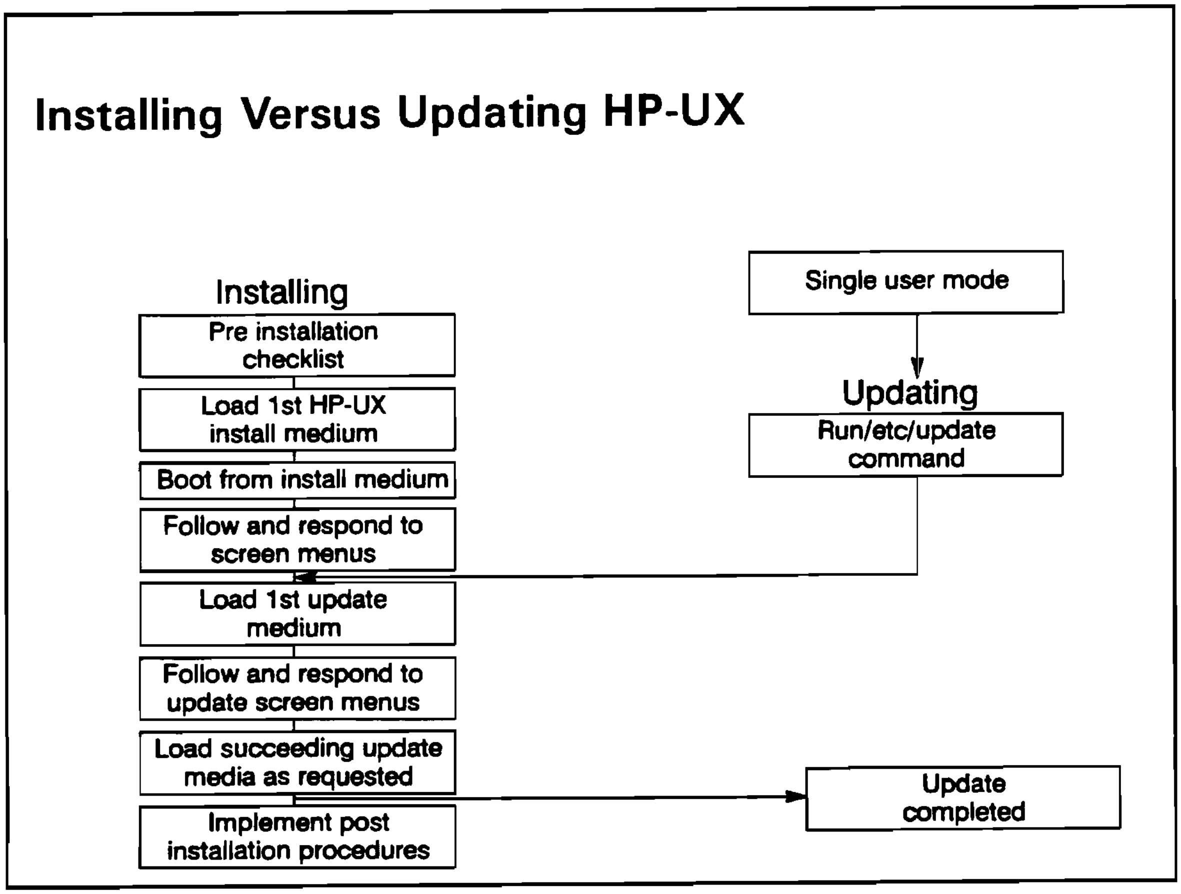

50722 B-1 129 ©1992 Hewlett-Packard Co.

Student Notes

The update process is the same as the "update" phase of the "install" process we covered in an earlier appendix. The relationship between the two processes is shown on the slide.

As a quick review of what we covered in the installation appendix: The install utility initialized the system disk, created a new file system, and copied over core HP-UX partitions and filesets and then rebooted the system. Then the install utility ran the `/etc/update` program to load the rest of the partitions and filesets to the system disk.

Remember that a **fileset** is a logical grouping of files and a **partition** is a logical grouping of filesets. A **product**, for example, the X Window System, is a logical grouping of filesets and/or partitions. When you load the filesets onto the disk from the update media, the files are copied to the disk in a hierarchical structure that is unrelated to the original fileset groupings on the update media.

When Hewlett-Packard releases a new version of the HP-UX operating system, such as version 8.0 or 9.0, you will use the update process to replace outdated portions of HP-UX with the current versions and add new features. When you update HP-UX, you load software products from the update media (or from a special network server) and incorporate that software (in the form of files and filesets) into an existing HP-UX file system. If any new or replacement files affect the existing kernel configuration, the kernel is reconstructed during the update. If the kernel is reconstructed, the update process will reboot your system.

Update Media Types

The type of update media you use depends upon your system type. Starting at the 8.0 release, new media types are supported on most HP 9000 systems. There are differences regarding supported update media types between the series 300/400, 700, and 600/800 families of computer systems.

Possible choices for media options are:

Table B-1. HP-UX 8.0 Media Options

Series 300/400	Series 700	Series 600/800
CD-ROM	CD-ROM	CD-ROM
cartridge tape	DDS/DAT	DDS/DAT
		mag tape
		cartridge tape

You may also update your system over a Local Area Network (LAN) provided that there is a computer system that has been configured as a network distribution (netdist) server. This method will be described in the end of this appendix. First we will talk about the update methods that use local physical media instead of network updates.

B-1. SLIDE: Installing Versus Updating HP-UX — Instructor Notes

Purpose

To review what an **update** is with respect to HP-UX. We covered this briefly in the installation appendix, but that was early in the course and only a few details were given. This appendix provides a more in-depth look at the **update** process.

Key Points

- Update is typically used when Hewlett-Packard releases a new version of the HP-UX operating system and the customer wants to utilize the new features of the update. Update can also be used to add a subsystem to HP-UX that was not added when the system was installed (or has been removed following installation).
- When you perform an update, you load products from supported update media, incorporate filesets into the file system, and rebuild the existing kernel (if necessary).
- Since this is a combined 300/400/700 and 600/800 appendix, we will have introduced media types that are not supported some system types.
- It is necessary to say a few things about the hardware components since we will be asking the student to insert media in the "drive" used for the update. The basic location of the update device may not have been covered earlier.
- CD-ROM media/drives will be first used at the 8.0 release for HP-UX software distribution for installation and update.
- The DAT/DDS tape drive is not supported at 8.0 for installing HP-UX and updating HP-UX on **ALL** systems. Although this is supported on the series 600/700/800 HP-UX systems it is NOT supported on series 300 and 400 at 8.0.
- Normally, you should use `update` interactively, which is what we will cover in this appendix. You might choose to use `update` non-interactively if:
  - You have a non-HP terminal that is incompatible with the screen control used in the interactive version of `update`.
  - You use `update` from a script or shell program.
  - You use `update` regularly and want to bypass interaction.

References

Since the HP-UX 7.0 release, the Series 300/400 and Series 600/800 `update` utilities have been merged. Starting at 7.0 all UPDATE media was distributed in `tar(1)` format. This remains the case at 8.0 (except for the addition of CD-ROM distribution media which has a unique format). Briefly, here's how the 8.0 update differs from previous HP-UX update programs:

- The interactive interface to `update` has changed in several ways:

- Some `update` screens have minor format changes.
- Some `update` screens have rewording of informational text.
- The main menu screen requires the entry of **codewords** for codeword protected distribution media (such as CD-ROM media)
- Update has a much better sense of fileset dependencies than before.
- Update uses a new tool, `rmfn` to reduce disk space if you need it during an update.
- The Series 300 and series 400 use cartridge tape and CD-ROM media for update media.
- The Series 600 and series 800 uses reel tapes, cartridge tapes, and DDS/DAT tapes for update media.
- The CD-ROM update runs unattended once you start loading. This eliminates the need for your being present to change tapes.
- The 8.0 `update` utility calculates time to complete, as it loads the software. This feature predicts the time remaining in the update process.
- The names and contents of filesets and partitions have changed dramatically in the 8.0 release. This affects the "Update only filesets on the system" option.

B-2. SLIDE: Before You Update HP-UX

Before You Update HP-UX

- Read the *HP-UX Release Notes* for the new release.
- Obtain codewords and hardware IDs if using codeword protected media.
- Make a note of the device file name for your installation device.
- Ensure that your TERM environment variable is set correctly.

50722 B-2 130

Student Notes

Read the *HP-UX Release Notes* for the new release.

This slide shows the steps you should take before you perform your update. You should carefully follow these steps, modifying them as directed by other manuals needed when you situation differs from a standard system (particularly if updating an HP-UX Clustered Environment).

Prior to updating the operating system, you should read the *HP-UX Release Notes*. Any new release will include numerous changes that may affect your users and applications running on your system. The *HP-UX Release Notes* is available in electronic form in the file `/etc/newconfig/8.0ReleaseNote` after you have completed an installation of the HP-UX 8.0 release. The *README FIRST* document describes how you can extract this file from your UPDATE media before you start the update process.

The *Software Status Bulletin* (SSB) provides information about currently outstanding problems and enhancement requests. You can get information, for example, about a known problem related to an HP-UX subsystem. An SSB will be included in the 8.0 system as the following file: `/usr/contrib/doc/SSB_8.0`.

You should carefully read the latest version of *Installing and Updating HP-UX HP 9000* before starting an update to 8.0.

Codewords for CD-ROM

Codewords are used with the `update` program to "unlock" software products stored on the CD-ROM disk in order to install or update these products on a specific computer system. Codewords are shipped to you from HP and are computed from information you supply with your order. This information includes the software product number that you have purchased and a unique hardware ID from the system you intend to install the software product on.

Device File Names

You must know the device file name for your installation device (CD-ROM or DDS/DAT drive, cartridge or reel tape drive) so that you know whether or not you need to override the default that `update` uses.

- On a Series 300/400/700, `update` uses `/dev/update.src` as its default.
- On a Series 600/800, `update` uses `/dev/rmt/0m` as its default.

If this is not the device file for the installation device you intend to use, you must change the source device name once you are in the `update` program.

Set TERM environment variable correctly.

If your `TERM` variable is not set properly, your display could behave strangely during an interactive update process.

B-2. SLIDE: Before You Update HP-UX — Instructor Notes

Purpose

To describe the steps you should take before updating your system.

Key Points

- Make a note of the name of the device file for your intended installation device. `update` uses the defaults `/dev/update.src` (Series 300/400/700) and `/dev/rmt/0m` (Series 600/800).
- If your `TERM` variable is not set properly, your display could behave strangely during an interactive update process.
- The appendix on installation should have explained the significance of codewords and hardware IDs when installing from CD-ROM media.
- Codewords will be required during the installation from CD-ROM.
- Codewords are derived from hardware IDs and are tied to specific software products.
- All current HP-UX software is actually on the CD-ROM disk, the codeword unlocks that product's filesets and partitions for installation (or update) on a specific system.
- Codewords are used only when using CD-ROM installation media. They may be applied to other distribution technologies in the future, but currently *only* CD-ROM media requires their use.
- All current HP-UX software products are stored on the CD-ROM disk that was shipped with your order. The software products are in effect "locked" and accessible only by codeword "keys". Subsequent purchase of other software products will not require media shipment, only a customer purchase order and a codeword delivery.

References

- *Installing and Updating HP-UX, HP 9000 Series 300/400* HP P/N B1862-90002 12/90
- *Installing and Updating HP-UX, HP 9000 Series 600/800* HP P/N B2437-90003 12/90
- *How HP-UX Works, Concepts for the System Administrator* HP P/N B1862-90005 12/90 (for all HP 9000 systems)
- *System Administration Tasks, HP 9000 Series 300/400* HP P/N B1862-90008 12/90
- *System Administration Tasks, HP 9000 Series 600/800* HP P/N B2437-90006 12/90

Transition

Let us consider this new security feature: codeword protected software distribution.

B-3. SLIDE: Codewords and Hardware IDs

Codewords and Hardware IDs

Codewords:

- Enable access to licensed software
- Computed based on hardware ID, software products purchased, and CD-ROM disk part number
- Supplied by HP in one of two forms:
  - Codeword Certificates for pre-computed codewords
  - Codeword Entitlement Certificates for post-computed codewords

Hardware IDs:

- Must be a unique identifying number related to your system hardware
- Are electronically readable by software
- Examples:
  - HP-HIL ID module
  - SW-ID number (Series 700)
  - Station (link level) address for LAN interfcae (Series 300/400)
  - HP-IB CD-ROM serial number

50722 B-3 131

Student Notes

All HP software distributed on CD-ROM comes loaded on the disks shipped to you as part of your HP-UX Software on CD Media subscription service. This service will provide two CD disks, one containing the operating system and subsystems and the other containing HP applications. Most of these products can only be accessed using a special codeword that unlocks the products. The codewords for these disks may be:

- Pre-computed codewords that are provided on Codeword Certificates shipped with the CD disks.
- Post-computed codewords; directions for obtaining post- computed codewords are provided on Codeword Entitlement Certificates.

Whether or not codewords are pre-computed depends on how the order was placed. One factor is the hardware ID. If the order is placed specifying the hardware ID as some hardware other than the HIL ID module, the codewords cannot be pre- computed.

B-3. SLIDE: Codewords and Hardware IDs — Instructor Notes

Purpose

To explain the significance of codewords and hardware IDs when (installing and) updating from CD-ROM media.

Key Points

Codewords will be required during (installation and) update.

Codewords are derived from hardware IDs and are tied to specific software products.

All current HP-UX software is actually on the CD-ROM disk, the codeword unlocks that product's filesets and partitions for installation (or update) on a specific system.

Codewords are used only when using CD-ROM installation media. They may be applied to other distribution technologies in the future, but currently ONLY CD-ROM media requires their use.

All current HP-UX software products are stored on the CD-ROM disk that was shipped with your order. The software products are in effect "locked" and accessible only by codeword "keys". Subsequent purchase of other software products will not require media shipment, only a customer purchase order and codeword delivery is needed.

Reference

Installing and Updating HP-UX (for HP 9000 series 300/400) HP P/N B1862-90002 12/90

B-4. SLIDE: Preparing the File System for Updating

Preparing the File System for Updating

- Make sure you have adequate disk space to contain the new release.
- Perform a full system backup.
- Clean up the HP-UX file system and check disks with `fsck`.
- Save a copy of your kernel to `/SYSBCKUP`.
- Ensure that any files that affect system configuration are current.

50722 B-4 132

Student Notes

Adequate Disk Space

If you are updating to HP-UX 8.0, you will *not* need more disk space then with the 7.0 release. Significant changes in HP-UX libraries have occurred that reduces the required space for HP-UX commands and HP supplied executables. These space savings were offset by increased functionality and new features, but not to the extent that more disk space is required. Of course if you are adding new applications, you must accommodate their space requirements.

Before you update your system, analyze the amount of space available on your destination disk. Make sure there is enough space to accommodate the update. Consider these factors:

- The files on your system.
- The size of your disk (or disks).
- The partitions and filesets you will 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.

The `update` program will tell you if you have insufficient disk space to update your system. However, if you suspect you might not have enough disk space to update your system, it is best to handle the problem before you begin. You can free disk space by removing files from your system or by creating symbolic links. You could also choose to mount another file system.

Perform a full system backup.

It is essential to perform a full system backup before you update your system. This will allow you to restore any or all of your files if something goes wrong with the update.

Clean up the HP-UX file system

You should start with a file system that is in good shape. You should have been periodically running `fsck` to check disks. It would be a good idea to do this before you start an `update` to be highly confident of the file system integrity.

Save a copy of your kernel to /SYSBCKUP

If the update process recreates the kernel, this step is not necessary, since it will be done automatically for you. If however, the `update` fails for some reason, it's a good idea to be able to reboot your system from a "known good" kernel.

In most cases this will not be a "bad" thing to do. There are some problems with simply assuming you can reboot from the saved kernel in /SYSBCKUP. For example, when updating to another major release of the HP-UX the backup kernel may be an earlier revision and therefore incompatible with the rest of the HP-UX files and commands.

Files that affect system configuration should be current

Ensure that any files that affect system configuration are properly backed up and correctly modified according to the installation instructions in your *README FIRST* documents. As we cover more about such files later in the course we will discuss various files that affect the system configuration, such as `/etc/checklist` and others.

B-4. SLIDE: Preparing the File System for Updating

Instructor Notes

Purpose

To describe the steps you should take before updating your system.

Key Points

- Even though `update` will tell you if you have insufficient disk space to update your system, you should check your disk space before performing an update. If you suspect you might not have enough disk space to update your system, it is best to add or delete space before performing the update.
- It is best to perform a full system backup before you update your system. This will allow you to restore any or all of your files if something goes wrong with the update.
- The `/etc/update` facility will mount any file systems listed in `/etc/checklist`. If you don't want a file system in this file to be mounted, comment out the line.
- Files that affect system configuration must be made current. For Series 600/800 users, ensure that your S800 file is properly backed up, along with a copy of `/etc/devices` These are to be used for future reference only, since at 8.0 the `/etc/devices` file will no longer be needed. The S800 file and the "/etc/devices" file may not have been covered in the course at this point, but we should mention this behavior here for later reference.

  The `/etc/update` facility will mount any file systems listed in `/etc/checklist`. If you don't want a file system in this file to be mounted, comment out the line. The mounting of file systems and the "/etc/checklist" file may not have been covered in the course at this point, but we should mention this behavior here for later reference.

B-5. SLIDE: Steps to Update HP-UX with Local Media

Steps to Update HP-UX with Local Media

1. Execute `shutdown` to bring the system to a single-user state.
2. Write-protect your *UPDATE* media.
3. Power-on the update device, insert media, and wait for drive to ready.
4. Copy update tools from media to disk.
5. Execute `/etc/update` and follow the interactive menus.

50722 B-5 133

Student Notes

This slide describes the steps to update you system software if you plan to use physical media such as tape or CD-ROM. If you plan to use a network server to update your system software across a network, then some steps will be different. The network update is covered separately in this appendix.

Execute `shutdown` to bring the system to a single-user state

Depending on the filesets that are being loaded, the system may or not be rebooted during the update process. Filesets that require kernel reconfiguration will reboot the system. If you unsure if you are updating this type of fileset, it is best to be prudent and bring your system to a single-user run-level with the `/etc/shutdown(1m)` command:

```
# cd /
# shutdown
```

We will talk more about shutting down the system later in the course. Do not perform this step if you intend to update across the network, since the command will terminate processes needed for LAN access.

Write protect your *UPDATE* media

If you are not using CD-ROM to update, the media you are using is vulnerable if not write protected. For cartridge tape media, turn the write-protect arrow towards the "Safe" label on the tape. For the DDS/DAT tape media, push the write-protect tab towards the corner of the tape. For mag tape media, remove the write-protect ring on the back of the tape reel, if present.

Power on the update device, insert media, and wait for drive to ready

Turn the power on for the update device. Then, after the device self test and warm-up are complete, insert the first UPDATE medium into the device. Now wait for the drive to become ready. Some devices will indicate readiness by turning on a light and others will turn off a light. Consult the hardware installation manual for your particular drive. Remember to put the device on "online".

Copy update tools from media to disk

The update tools, the `update` program and its supporting files and programs, are written to support a particular kernel revision. Therefore, if updating for the previous version of the operating system to the current, the current update tools must be extracted from the media and downloaded to disk.

Execute `/etc/update` and follow the interactive menus

Make sure you are in the root (/) directory before invoking `update`.

You should type `/etc/update` on the HP-UX command line to invoke interactive update process. The menus will lead you through the process of updating the system.

B-5. SLIDE: Steps to Update HP-UX with Local Media

Instructor Notes

Purpose

To describe the steps necessary to update your system. We continue the details on the next slide.

Key Points

- It is best to shutdown your system to single-user mode before updating. Do not do this for a network based update.
- Make sure your writable UPDATE media are write protected before inserting them into the update device.
- You must copy the new `update` program for the tape to the disk before invoking `update` (This is usually true for a major release such as 8.0, see the README FIRST document to be sure). Details on the next slide for this important step.
- After you invoke `update`, the program is interactive and will lead you through the necessary steps to update your system.

Transition

We continue to give more details about the process on the next slide.

B-6. SLIDE: Copying the Update Tools

Copying the Update Tools

To copy the /etc/update program and other tools:

- for Cartridge tape media

```
# cd /
# tcio -iZ /dev/dev-file-name | tar -xvf - TOOL
```

- for DDS/DAT and mag tape media

```
# cd /
# tar -xvf /dev/dev-file-name TOOL
```

- for CD-ROM media

```
# cd /
# /etc/mount /dev/bsrc /UPDATE_CDROM  -r  -t cdfs
# tar -xvf /UPDATE_CDROM/TOOL
```

50722 B-6 134

Student Notes

There are different commands to copy the `update` program and other required tools from the update media. The command you use depends upon the type of media you are planning to use. In the commands given on the slide, you will need to supply the correct device file name for your update device. The device file name you use depends upon your particular system configuration. The default device files are usually present on your system, having been put there by the install process. The device file name you might use is:

- `/dev/update.src` for cartridge tapes.
- `/dev/update.src` for DAT tape.
- `/dev/rmt/0m` for reel tape.

Notice that the CD-ROM command is a different format than for other media types. Usually the location where the CD-ROM is mounted is /UPDATE_CDROM, but the location may vary on your system. If the CD-ROM drive is not mounted, it must first be mounted before you type this command.

After you execute the `tar` command to copy the `update` program on to your disk, you will have to wait several minutes for the first TOOL file to be extracted. Once the process begins, the files are extracted rapidly. The TOOL files are located at the beginning of the tape; however, `tar` is unaware of their location and as such will continue to search for these files until the end of the tape. To shorten this process, when the last TOOL file is echoed to display, wait 5-7 minutes to be sure it actually is the last TOOL file, and then press (Break) to end the `tar` process.

B-6. SLIDE: Copying the Update Tools — Instructor Notes

Purpose

To describe in more detail the steps necessary to copy the update system from the update media.

Key Points

You must copy the new `update` program for the tape to the disk before invoking `update` before every update. This is mandatory since the `update` program changes with new releases, and the previous `update` command (in the last release) may not match the format of the data on the Update media. The version you copy (in this step) from the media is guaranteed to match the format on the media.

- Directories updated by getting the TOOL fileset from the update media are: `/system` `/etc/update.lib` `/etc/interface.lib` `/etc/rmfn.lib`
- System commands updated by getting the TOOL fileset from the update media are: `/etc/update` `/etc/updist` `/etc/netdistd` `/etc/rmfn`

References

Note that the `tar` *key* argument can be preceded by a hyphen (-), but it is not necessary. For example:

```
tar -xvf . . .
    is equivalent to
tar xvf . . .
```

On careful inspection of the list of files extracted by the `tar(1)` command, you will discover that the `sysrm(1m)` command is no longer part of the TOOL fileset. It has been replaced by the `rmfn(1m)` command.

The following example shows the command and output for a Series 300/400 system.

```
# tcio -iZ /dev/update.src | tar -xvf - TOOL

x system/INDEX, 172 bytes, 1 tape blocks
x system/INFO, 1799 bytes, 4 tape blocks
x system/CDFinfo, 395 bytes, 1 tape blocks
x TOOL/../etc/rmfn, 299008 bytes, 584 tape blocks
x TOOL/../etc/netdistd, 217088 bytes, 424 tape blocks
x TOOL/../etc/update, 536576 bytes, 1048 tape blocks
TOOL/../etc/updist linked to TOOL/../etc/update
x TOOL/../etc/rmfn.lib/data_f.dic, 2049 bytes, 5 tape blocks
x TOOL/../etc/rmfn.lib/screen_f.lib, 11086 bytes, 22 tape
blocks
x TOOL/../etc/update.lib/data.dic, 909 bytes, 2 tape blocks
```

```
x TOOL/../etc/update.lib/global.ini, 71 bytes, 1 tape blocks
x TOOL/../etc/update.lib/screen.lib, 59983 bytes, 118 tape
blocks
x TOOL/../system/TOOL/CDFinfo, 394 bytes, 1 tape blocks
x TOOL/../system/TOOL/copyright, 461 bytes, 1 tape blocks
x TOOL/../system/TOOL/customize, 249 bytes, 1 tape blocks
x TOOL/../system/TOOL/decustomize, 321 bytes, 1 tape blocks
x TOOL/../system/TOOL/kern_util, 34776 bytes, 68 tape blocks
x TOOL/../system/TOOL/rebuild.ckerns, 21205 bytes, 42 tape
blocks
x TOOL/../system/TOOL/kernrel, 11 bytes, 1 tape blocks
x TOOL/../system/TOOL/pdf, 1569 bytes, 4 tape blocks
x TOOL/../etc/interface.lib/jam/hpitevid.bin, 1052 bytes, 3
tape blocks
x TOOL/../etc/interface.lib/jam/hpkeys.bin, 238 bytes, 1
tape blocks
x TOOL/../etc/interface.lib/jam/hpvid.bin, 1027 bytes, 3
tape blocks
x TOOL/../etc/interface.lib/jam/msgfile.bin, 4125 bytes, 9
tape blocks
x TOOL/../etc/interface.lib/jam/vt100keys.bin, 422 bytes, 1
tape blocks
x TOOL/../etc/interface.lib/jam/vt100vid.bin, 850 bytes, 2
tape blocks
x TOOL/../etc/interface.lib/jam/wy30keys.bin, 230 bytes, 1
tape blocks
x TOOL/../etc/interface.lib/jam/wy30vid.bin, 842 bytes, 2
tape blocks
```

For comparison, the previous release extracted the following files from the TOOL fileset:

```
x TOOL/../etc/sysrm, 52420 bytes, 103 tape blocks
x TOOL/../etc/netdistd, 204800 bytes, 400 tape blocks
x TOOL/../etc/update.6.5, 139264 bytes, 272 tape blocks
x TOOL/../etc/update, 438272 bytes, 856 tape blocks
x TOOL/../etc/updist linked to TOOL/../etc/update
x TOOL/../etc/update_scr/data.dic, 677 bytes, 2 tape blocks
x TOOL/../etc/update_scr/global.ini, 54 bytes, 1 tape block
x TOOL/../etc/update_scr/hpitevid.bin, 1050 bytes, 3 tape
blocks
x TOOL/../etc/update_scr/hpkeys.bin, 238 bytes, 1 tape block
x TOOL/../etc/update_scr/hpvid.bin, 1025 bytes, 3 tape
blocks
x TOOL/../etc/update_scr/msgfile.bin, 4167 bytes, 9 tape
blocks
x TOOL/../etc/update_scr/updlib, 49562 bytes, 97 tape blocks
x TOOL/../system/TOOL/CDFinfo, 558 bytes, 2 tape blocks
x TOOL/../system/TOOL/copyright, 455 bytes, 1 tape block
x TOOL/../system/TOOL/customize, 760 bytes, 2 tape blocks
x TOOL/../system/TOOL/kern_util, 35156 bytes, 69 tape blocks
```

```
x TOOL/../system/TOOL/rebuild.ckerns, 14106 bytes, 28 tape
blocks
x TOOL/../system/TOOL/kernel, 11 bytes, 1 tape blocks
x TOOL/../system/TOOL/pif, 2982 bytes, 6 tape blocks
x TOOL/../system/TOOL/revlist, 597 bytes, 2 tape blocks
```

Transition

We will now look at the first screen presented to us after the `update` program starts running.

B-7. SLIDE: Update Main Menu

Update Main Menu

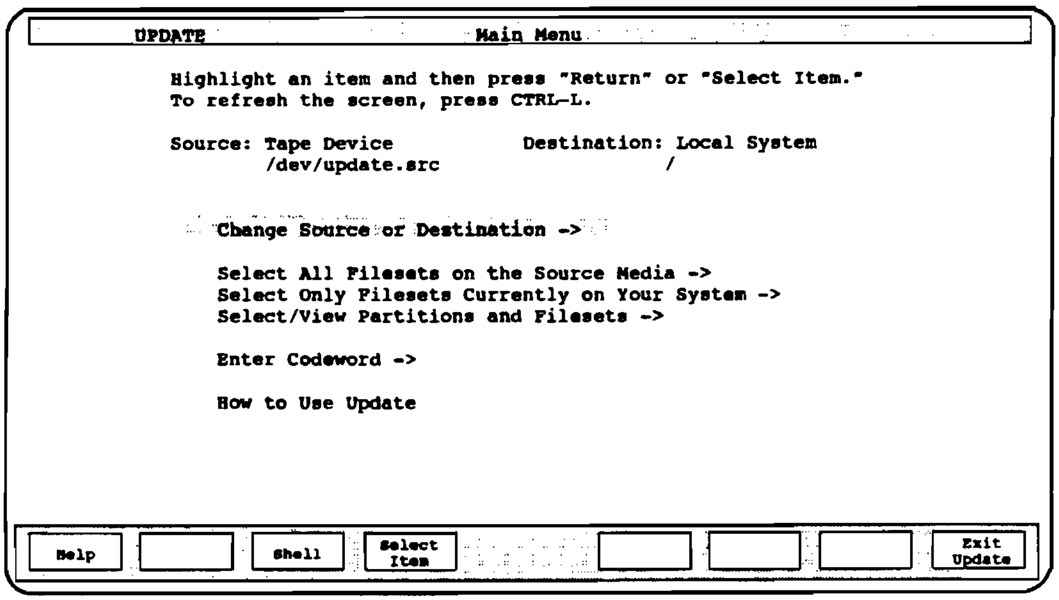

50722 B-7 135

Student Notes

Navigating the Screen

Once you invoke `/etc/update`, you will see a message:

```
Initializing ...
```

followed by the screen as seen on the slide: Update Main Menu. You navigate in the `update` program the same way you navigate in SAM. To select an item:

- Use (Tab) or the cursor (arrow) keys to highlight the option you want to select.
- Press (Return) or `Select Item` to choose the current selection.

- Press (Tab) or (Shift)(Tab) to move forwards and backwards between the fields on the screen if a form is displayed..

If you are unsure about the function of certain menu items or data entry fields, highlight the item and then press `Help` for on-line information about the item. This capability applies only to those "update" screens where a softkey labeled `Help` appears.

If you want to quit the `update` process before going any further, press `Exit Update`. This capability applies only to those "update" screens where a softkey labeled `Exit Update` appears.

Changing the Source or Destination

Select "Change Source or Destination" if your source or destination differs from the default listed on the Main Menu. A new menu will appear that allows you to enter the correct source and/or destination device. A sample of that menu (which is a pop-up type) follows in the next slide.

You would change the destination only if directed by the *README FIRST* notices shipped with your update media.

You may see several variations of the source specification:

```
Source: Tape Device          Destination: Local System
        /dev/update.src                   /

Source: Tape Device          Destination: Local System
        /dev/rmt/0m                       /

Source: Tape Device          Destination: Local System
        /UPDATE_CDROM                     /
```

Which Filesets and Partitions To Select?

You should choose to "Select All Filesets From the Source Media." This ensures that you get all the files necessary for a proper update. If you decide you do not need some of the filesets after your update, you can delete a file set with the `rmfn` (remove function) command. This new command supplants the `sysrm` command.

If you choose to "Select/View Partitions and Filesets," you will get a menu listing all the possible partitions that you can load. You can use it to see which filesets belong with which partition. This menu is shown in a later slide.

Be aware that the contents of filesets and partitions have been rearranged from HP-UX 7.0. The names of partitions and filesets have also changed for the 8.0 release. Therefore, if you choose to "Select Only Filesets Currently on your System," there is a possibility that all files might not be updated.

B-7. SLIDE: Update Main Menu — Instructor Notes

Purpose

To explain the basics of the Main Menu of the `update` program.

Key Points

- The "Source" default values change depending on the hardware system type used. The defaults differ for the series 300/400 or series 600/800 `update` programs.
- Both versions of the slides are created for you.
- If you call attention to the "Enter Codeword" menu choice, you should only do so for the Series 300/400 version of the course. The CD-ROM media that requires this "feature" is only supported (at 8.0 release) for `update` on the series 300/400 systems. Later HP-UX releases may support it on the Series 600/800 systems.
- You may have to change the source or destination depending on your system and choice of installation media.
- You will usually choose to "Select All Filesets on the Source Media." This ensures you get all the files necessary for a proper update.
- You can remove filesets that you don't want LATER with the `rmfn(1m)` command.
- The `rmfn(1m)` command replaced the `sysrm(1m)` command.

  Fileset definitions and names have changed for the 7.0 release. If you choose to "Select Only Filesets Currently on Your System" you should be careful to verify that all the filesets were correctly updated.

Note

Be aware that the contents of filesets (and partitions) have been rearranged. The names of partitions and filesets have also changed for the 8.0 release. A given fileset table of contents may differ between 8.0 and previous releases. You should verify that all obsolete files were removed following an update. Also, if you choose to "Select Only Filesets Currently on your System," there is a possibility that all files might not be updated.

Screen Changes from 7.0

You may notice some changes relative to the HP-UX 7.0 version of the `update` Main Menu. There are several format changes and a different ordering of the menu selections.

The label: "Load Everything from Source Media" has changed to: "Select All Filesets on the Source Media"

The label: "Update Only the filesets on the Destination" has changed to: "Select Only Filesets Currently on your System"

The label: "View or Select Individual Partitions" has changed to: "Select/View Partitions and Filesets"

The choice "Change Source or Destination" has been moved to the top of the list to become the default selection.

Transition

This is the main menu screen for update. We may have to change the source or destination before we start selecting filesets to update.

B-8. SLIDE: Changing the Source or Destination

Changing the Source or Destination

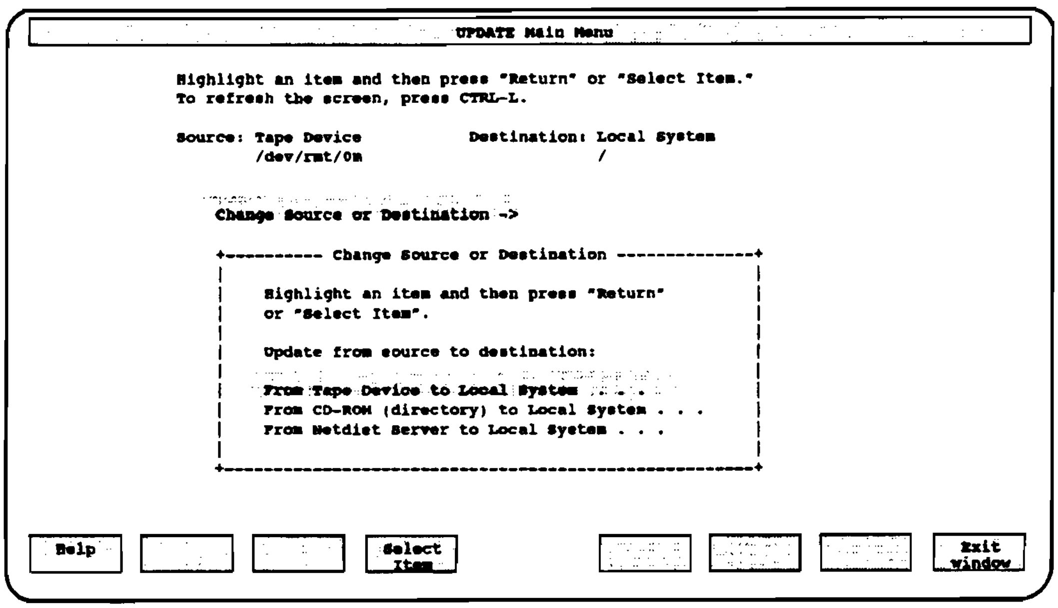

50722 B-8 136 ©1992 Hewlett-Packard Co.

Student Notes

If you select the update Main Menu choice: "Change Source or Destination" a **pop up** menu window will appear over the Main Menu window. This is shown in the slide above.

If you are updating from a CD-ROM drive, you need to change the source device by selecting `From CD-ROM (directory) to Local System`. After selecting the menu choice, the source will be shown on the Update Main Menu as:

```
Source: Tape Device          Destination: Local System
        /UPDATE_CDROM                     /
```

The CD-ROM needs to be mounted for the update to be successful. Since the update tools are extracted from the CD-ROM, it may very well be already mounted. If you did not mount the required CD-ROM in the installation drive to the directory /UPDATE_CDROM before starting the update, you will not be able to do so inside the update program from the shell. You must exit the update program to mount the correct CD-ROM file system and then restart update.

```
# /etc/mount /dev/bsrc /UPDATE_CDROM -r -t cdfs
```

If you are updating from a remote system using the LAN, you need to change the source by selecting From Netdist Server to Local System instead of the default. We will cover "netdist" updates in separate slide near the end of this appendix.

B-8. SLIDE: Changing the Source or Destination Instructor Notes

Purpose

To show the reaction to choosing the main menu selection: "Change Source or Destination". The Main menu is overlaid with another window in the form of a pop-up menu. After choices are made in this new window the window will disappear, hence the term: pop-up.

Key points

- The menu choices that end in " ... " will require further input from the user. This is usually an interactive dialog of some type, often a "fill in the blanks" type of pop-up form that must be completed.
- The use of CD-ROM update media requires mounting the `cdfs` file system to the root file system BEFORE running `update`. Failure to do so causes a fatal error.

Teaching Tips

You cannot skip teaching about the necessity for mounting a CD-ROM file system before updating. If the `mount(1m)` command was not covered before now, tell the students that we will cover it in more detail later in the class.

Transition

If you make the pop-up menu selection: `From CD-ROM (directory) to Local System`, you will see another pop-up screen that requires input.

B-9. SLIDE: From CD-ROM to Local System

From CD-ROM to Local System

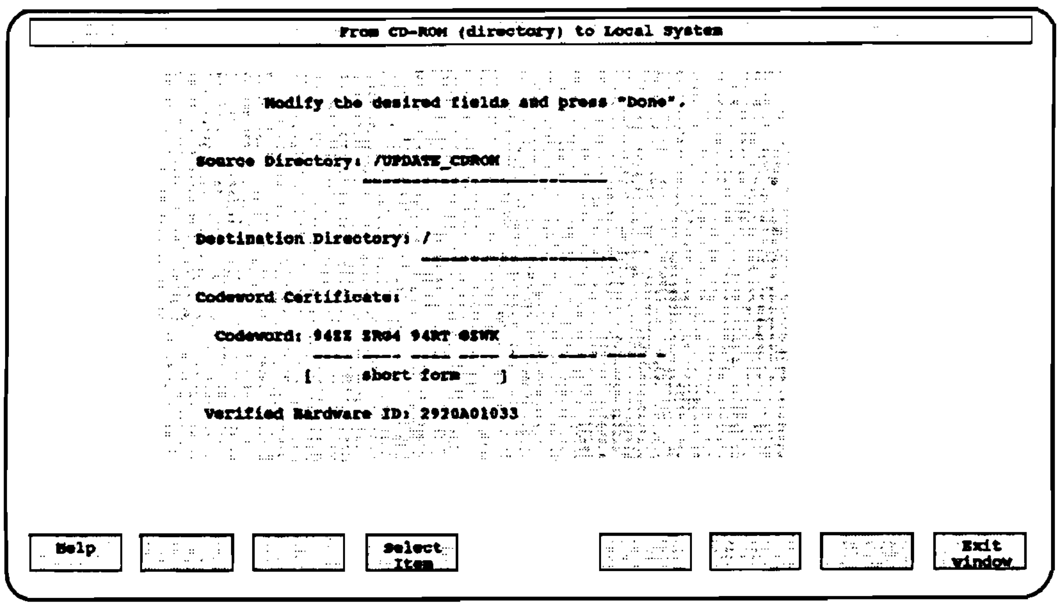

50722 B-9 137

Student Notes

When updating from codeword protected media such as CD-ROM disks, you must enter a codeword before selecting filesets. The codeword for your system software can be entered in the screen displayed here. The codeword you enter unlocks certain filesets on the "UPDATE" CD-ROM disk so that only the software products you purchased for this hardware can be installed.

Tab down to the codeword field and enter either the 16 or the 29 character codeword you have been issued for this hardware. The codeword can be entered in either upper or lower case for letters. Do *not* enter spaces that may appear for legibility in the codeword on your software product certificate.

The code words are tied to hardware IDs. Possible choices for hardware IDs are:

- A machine ID number such as the SW_ID field in Series 700 Stable Storage,
- A machine ID number from the LAN link level address on a Series 300/400 system,
- A security number from an HP HIL 46084A ID module on a Series 300/400/700.
- A security number from an HP C1707A HP-IB CD-ROM drive on a Series 300/400.

You should not enter anything in the field "Verified Hardware ID." This field will be filled in by the `update` program if the codeword you supply matches one of the possible hardware IDs on your system. If the field is filled in and the window pops down before you can read it, just choose that option again. In the second go around, don't change anything, just select **`Exit Window`** to return to the "Main Menu" screen.

B-9. SLIDE: From CD-ROM to Local System — Instructor Notes

Purpose

To explain the pop-up dialog of the `update` programs `From CD-ROM (directory) to Local System` selection inside the menu allowing a change of Source or Destination prior to stating the loading of files.

Key Points

- The codeword entry portion only applies to codeword protected software distribution formats. Currently this only applies to CD- ROM disc.
- A codeword entered here is stored in the file `/etc/update.lib/codeword` for later use on subsequent updates.
- Entering the codeword is essential if the update is to complete correctly. Files may not be loaded without a proper codeword.
- Review the sources of possible choices for hardware IDs used to generate codewords.
- The field "Verified Hardware ID" will be filled in by the `update` program.

Transition

If you make the pop-up menu selection: `From Tape Device to Local System`, you will see another pop-up screen that requires input.

B-10. SLIDE: From Tape Device to Local System

From Tape Device to Local System

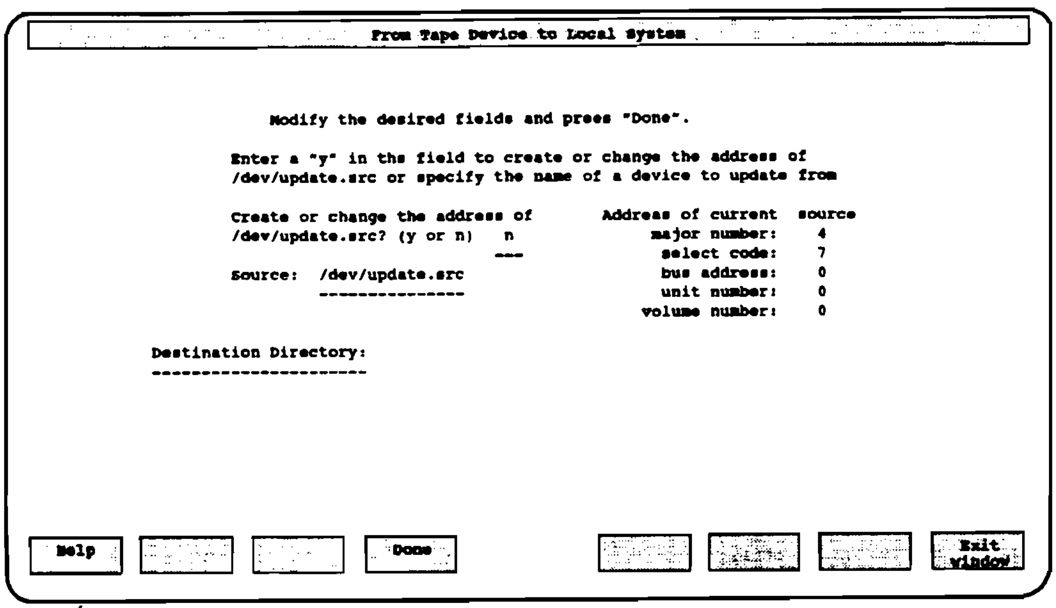

50722 B-10 138

Student Notes

When updating from tape media, you are not required to enter a codeword.

Since you got here by choosing `From Tape Device to Local System`, you can interact with this screen to change the source or destination from the defaults displayed here. You do not have to enter anything. You complete the interaction with `Done` or `Exit Window`. `Exit Window` allows you to exit without making any changes even after you modified fields on the screen.

B-10. SLIDE: From Tape Device to Local System Instructor Notes

Purpose

To explain the pop-up dialog of the update programs From Tape Device to Local System selection inside the menu allowing a change of source or destination prior to stating the loading of files.

Key Points

- When updating from tape media, you are not required to enter a codeword.
- The default source device may be different for different system types, as described earlier. No duplicate slide for series 300/400 and series 600/800 is provided.
- Review the device files if necessary at this point.

Transition

If you make the pop-up menu selection: From Netdist Server to Local System, you will see another pop-up screen that requires input.

B-11. SLIDE: From Netdist Server to Local System

From Netdist Server to Local System

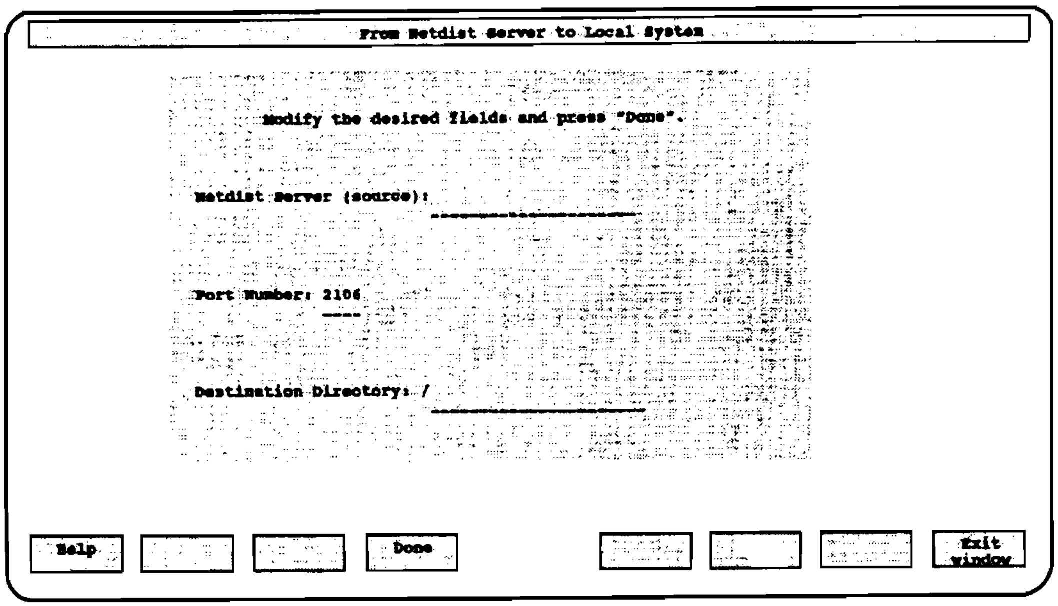

50722 B-11 139

Student Notes

Netdist Server Selection

When updating from a netdist server, you may be required to enter a codeword.

Since you got here by choosing `From Netdist Server to Local System`, you must interact with this screen to identify the netdist server or to change destination from the default displayed on the screen. You must name a host that has been properly configured to act as a netdist server. (We will cover configuration of netdist servers later).

You complete the interaction with `Done` or `Exit Window`. `Exit Window` allows you to exit without making any changes even after you modified fields on the screen.

B-11. SLIDE: From Netdist Server to Local System

Instructor Notes

Purpose

To explain the pop-up dialog of the `update` programs `From Netdist Server to Local System` selection inside the menu allowing a change of Destination and to name the remote network software distribution server prior to stating the loading of files.

Key Points

- We will cover the setup of the remote netdist server system later in this appendix.
- The default source system must be specified.
- The TCP Port number differs depending on the revision level of HP-UX on the target.
- The port number you use depends upon the revision level of the destination system. If a 7.0 system is running `update`, then the port number 2106 is correct. If an 8.0 system is being updated the port should be 2107.
- The port number you use depends upon the revision level of the destination system. If a 7.0 system is running `update`, then the port number 2106 is correct. If an 8.0 system is being updated the port should be 2107.

Transition

After setting the Source and Destination correctly, using the Pop-up window, we returned to the Main Menu. In the main menu we select one of the three options to update filesets on the destination system. The next slide will show us what filesets will be updated before we actually start loading files. Assume that we selected the option: "Select/View Partitions and Filesets". This is the screen we will see:

B-12. SLIDE: Update Individual Partitions Screen

Update Individual Partitions Screen

UPDATE View or Select Individual Partition

Source: /UPDATE_CDROM   Destination: /

Selected	Name	Partition Description	Size in Kbytes
n	DATABASE	Database Products	22724
n	GRAPHICS	Graphics Products	12071
n	NETWORKING	Networking Products	9891
n	NLIO	Native Language I/O Products	26757
n	NLS	Native Language Support	592
n	OS-ADMIN	Recommended Administration Cmds	7980
n	OS-CORE	Recommended System Core	11650
n	OS-FEATURES	Selectable OS Features	18626
n	PROG-LANGUAGES	Programming Languages	10564
n	REFERENCE-DOC	Reference Manual Pages	5813
n	SHARED-LIBS	Runtime Shared Libraries	4322
n	WINDOWS	Windowing Products	22703
v	vvvvvvvvvvvvvvv	vvvvvvvvvvvvvvvvvvvvvvvvvvvvvvvv	vvvvvvvvvvvvv

Help | | Shell | Start Loading | Disk Space | View Filesets | Global Select | Main Menu

Student Notes

If you choose to "Select/View Partitions and Filesets" from the Main Menu, then you will get this menu. A "working" message may appear at the bottom of the Main Menu screen for some time before this screen is displayed.

The time for the "working" message varies from a minute on some systems for a local update, but may range up to several minutes on remote (across the LAN) updates. You should be patient, a lot of information about the files you will be updating is being gathered from the update media or the netdist server.

Above the Source and Destination, you will see a message that says:

```
Mark "y" or "n".  To pick and choose individual filesets within a partition,
press "View Filesets".  A "p" means that some filesets have been selected
within a partition.  Press "Start Loading" when selection is complete.
```

If you mark "y" next to a partition, then you have chosen all the filesets within that partition. If you only want to choose some of the filesets in a partition, then highlight the partition and press `View Filesets`. This will take you to a new menu that lists the filesets for that partition. Mark the individual filesets that you want with a "y" and then press `Done`. When you get back to the Partition Menu, you will see a "p" (for partial) next to the partition.

The 8.0 (or later) version of `update` has much more knowledge of fileset dependencies than earlier versions. You may get warning messages if you select a fileset and then decide to "de-select" it. The warnings simply mean that you attempted to remove files that are needed by other selected filesets (there are hidden dependencies). This is actually a great help, but periodic warnings as you pick and choose filesets may be a bit distracting. The `update` program will remember your attempts to delete files and filesets it has warned you about. If you then remove the dependencies, the previous files will be automatically marked for removal. This is a bit of a surprise the first time it happens, but will prevent the need to circle back through the list multiple times.

In general you are much safer by loading everything (if you have disk space) and then using the `rmfn(1m)` command to delete software functions/products you do not want. Otherwise you will need to know which filesets to not load during the update.

Once you have marked all the partitions (and filesets) that you want to load, press `Start Loading`. Once you do this you will get several messages from the `update` program.

Next, you will get a "Calculating disk space requirements" message that flashes at the bottom of the screen. `update` will warn you if available disk space is insufficient. If this happens, a new screen is available that allows you to delete either some of the currently selected files you plan to update or existing disk files without leaving `update`. This will happen before you begin loading.

Messages

In the case of a single update medium, such as CD-ROM or netdist updates, you will get one last chance to confirm your intentions, or exit from the update process before loading begins. You will see a message such as:

```
The update should complete without additional attention.

Will load 52 filesets (5813 Kbytes), including 52 filesets (5813 Kbytes)
directly selected and 0 filesets (0 Kbytes) selected due to dependencies.
Review the log file, /tmp/update.log, afterward by:

1) typing "more /tmp/update.log";

2) finding the correct date and time heading;
```

```
3) looking for any messages that begin with ERROR or WARNING.


This is your last chance to change your mind about your selections.
Continue?  (y or n)
```

In the case of multiple media (tapes) you might get a message similar to the following:

```
 NOTE:  Filesets are selected that reside on multiple media.  You must
        change media at some point during the update.
In spite of these warnings or notes, you might want to continue.


Will load 79 filesets (25813 Kbytes), including 52 filesets (21403 Kbytes)
directly selected and 27 filesets (4410 Kbytes) selected due to
dependencies.


Review the log file, /tmp/update.log, afterward by:

1) typing "more /tmp/update.log";

2) finding the correct date and time heading;

3) looking for any messages that begin with ERROR or WARNING.


This is your last chance to change your mind about your selections.
Continue?  (y or n)
```

This means that you will need to change tapes during the update process. Once you choose to continue, `update` will start loading filesets from the first update tape to your disk. The next slide shows the screen you will get while `update` is loading files.

Note The `update` program might reboot the system as part of the update process. The program reboots the system if you've selected 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 have the option of proceeding or of exiting the `update` program.

B-12. SLIDE: Update Individual Partitions Screen   Instructor Notes

Purpose

To show the students what menu they will get if they choose to "View or Select Individual Partitions."

Key Points

Review filesets and partitions. This was covered in the Installation appendix, but now is a good time to review it.

Transition

Next we will look at the screen when we select "Begin Loading", entitled: "Loading Partitions and Filesets".

B-13. SLIDE: Loading Partitions and Filesets (Multiple Media)

Loading Partitions and Filesets (Multiple Media)

```
UPDATE                  Loading Partitions and Filesets

      Loading fileset:  ALLBASE-MAN       From /dev/update.src

                                  From medium number:  1 of 3

      This Medium:  Kbytes loaded:            8106   of     51059
    (4%)
                    Loading fileset:             3   of        36

      All Media:    Kbytes loaded:            8106   of    117072
                    Loading fileset:             3   of        97

              Estimated remaining hours:minutes:              1:43

          Summary of Messages (also logged to /tmp/update.log)

              -- searching for next fileset --

 Help | Main Menu | Shell | Perform Task |  |  | Disk Info | File Sys Info | Exit Task
```

50722 B-13 141

Student Notes

This is the screen you will get while `update` is loading your files from the update media to the disk. It summarizes the activity of the `update` program by telling you which medium you are loading, how many kilobytes have been loaded, and which fileset it is on. When the program goes from one fileset to the next, you will see the message:

```
-- searching for next fileset --
```

This lets you know that it is looking for the next fileset on the update media.

Once all the filesets are loaded, your update is complete. If necessary, the system will reboot. Log in and look at the update logs.

B-13. SLIDE: Loading Partitions and Filesets (Multiple Media)

Instructor Notes

Purpose

To show the screen that comes up while `update` is actually loading partitions and filesets.

Key Points

After all the filesets are loaded, your update is complete. If necessary the system will reboot. The first thing you should do, is look at your update logs.

Transition

The next slide show the screen we see if a CD-ROM update or a netdist update is being used. There is only one medium in these cases, so the screen is simplified:

B-14. SLIDE: Loading Partitions and Filesets (Single Media)

Loading Partitions and Filesets (Single Media)

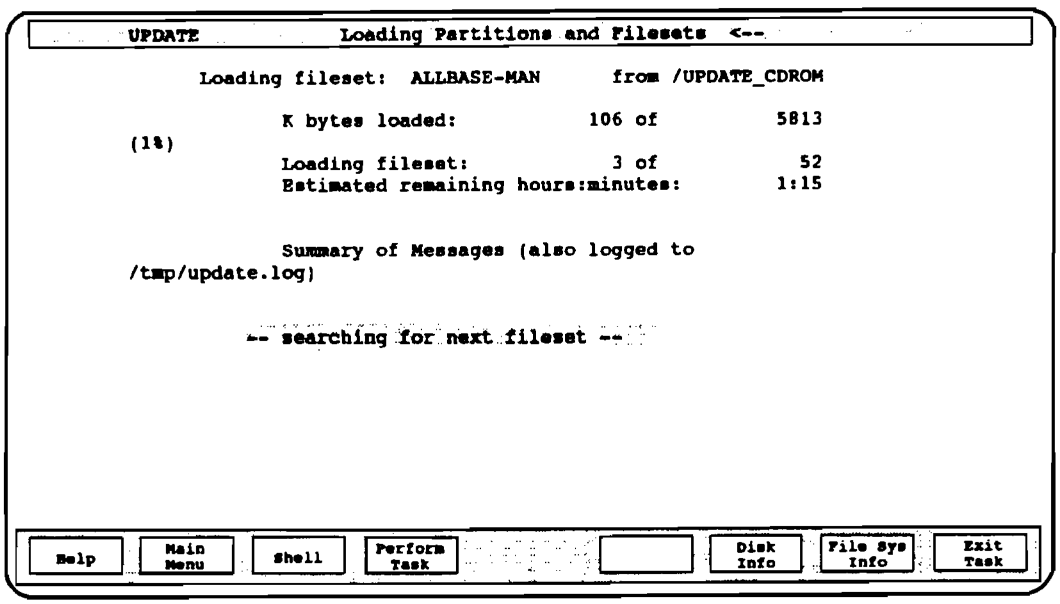

50722 B-14 142

Student Notes

This is the screen you will get while `update` is loading your files from the update medium to the disk. It summarizes the activity of the `update` program by telling you which medium you are loading, how many kilobytes have been loaded, and which fileset it is on. When the program goes from one fileset to the next, you will see the following message:

```
-- searching for next fileset --
```

Usually this message will flash at the bottom of the screen. The message lets you know that the `update` program is searching for the next fileset on the update medium.

Once all the filesets are loaded, your update is complete. If necessary, the system will reboot. Log in as normal and look at the update logs.

B-14. SLIDE: Loading Partitions and Filesets (Single Media)

Instructor Notes

Purpose

To show the screen that comes up while `update` is actually loading partitions and filesets, if a single update medium is being used. This is the case with a network based update or CD-ROM based update.

Key Points

After all the filesets are loaded, your update is complete. If necessary the system will reboot. The first thing you should do, is look at your update logs.

Transition

The next slide show what actions and where to look after an update.

B-15. SLIDE: Summary of the Update Process

Summary of the Update Process

- In Update Main Menu, select "Change Source or Destination"
- In "Change Source or Destination" Menu, select one of:
  - "From Tape Device to Local System"
  - "From CD-ROM (directory) to Local System"
  - "From Netdist Server to Local System"
- Interact with pop-up windows, giving required information
- Return to the Update Main Menu and select one of:
  - "Select All Filesets on the Source Media"
  - "Select Only Filesets currently on your System"
  - "Select/View Partitions and Filesets"
- Mark the desired filesets and partitions to load
- Press `Start Loading`
- Wait for completion, change media as required

50722 B-15 143

Student Notes

B-15. SLIDE: Summary of the Update Process

Instructor Notes

Purpose

To summarize the coverage of this appendix, so far.

Key Points

This slide covers the basic steps covered so far in the Update Process.

Some steps differ slightly based on your system type.

Transition

The next slide show what actions and where to look after an update.

B-16. SLIDE: After the Update

After the Update

- Check the /tmp/update.log file
- Check the /etc/newconfig directory
- Make a backup kernel (in some cases)
- Make a full backup
- Make a recovery system (Series 300/400/700 only)

50722 B-16 144

Student Notes

The /tmp/update.log

The /tmp/update.log file contains a description of the events and any errors that occurred during the update process. The following items are message labels and their meanings. The actual messages are indented nine spaces.

The label ======= indicates that a task within update is beginning or has completed. For example:

```
=======...BEGINNING UPDATE PROGRAM
```

The label ERROR: indicates that the program cannot proceed, or that it needs corrective action. In some cases this impacts update so much that it cannot continue. An example of a severe error is:

```
ERROR:  Destination directory "/mount" is invalid:  No such file or directory.
```

The label `WARNING:` usually indicates the program can continue. However, something went wrong or requires attention, either now or later. One example that can often be ignored is:

```
WARNING: Cannot access /etc/checklist file: No such file or directory.
```

Another example of a warning that is more serious and may prevent `update` from working as expected is:

```
WARNING: Cannot find any sources of hardware IDs for codeword verification.
         See the System Administrator Manual.
```

The `NOTE:` label indicates that something out of the ordinary or worth special attention has happened. The message may require no action on your part, in other cases the `NOTE:` message will require action. In some cases you must infer the action you must take after the update. For example:

```
       * Beginning to load fileset "TOOL".
NOTE:    Did not remove file "/etc/#update":  Text file busy (errno = 26).
       * Successfully loaded fileset "TOOL".
       * Beginning to load fileset "UX-CORE".
NOTE:    Did not remove file "/etc/#init":  Text file busy (errno = 26).
       * Successfully loaded fileset "UX-CORE".
```

You will have to remove the duplicate copy of the `update` and `init` programs which have been renamed: `/etc/#update` and `/etc/#init` after the update completes.

Checking the /etc/newconfig Directory

The `/etc/newconfig` directory contains new versions of some system files normally put in the `/etc` directory (for example, `rc`, `brc`, `backup`). Because you might have edited the original versions of these files, they are not replaced by either `update` or its customized scripts. Use the `diff` command to find the differences between your old files and new files placed in `/etc/newconfig`. Then you can incorporate your changes into the new versions of these files and move the new files into the `/etc` directory. Check the `README` file for information about the files in `/etc/newconfig`.

If you have problems with your update see the chapter, Updating HP-UX, of *Installing and Updating HP-UX* for your system type. Another good place to look is *Solving HP-UX Problems*, a new book at the 8.0 release.

Make a Backup Copy of the Kernel

You only need to make a new backup kernel under certain conditions:

- If updating from one major revision of the operating system to another, as in 7.0 to 8.0. In that case the kernel in /SYSBCKUP may not work with your new system files.
- The update did not make a backup kernel for you.

You can tell if you need to do backup the kernel by reading the update.log file described above. The update.log will contain messages to this effect if the kernel was backed up. A Series 300/400/700 example of these messages in the log file follows:

```
          ...
NOTE:     Copied /etc/conf/dfile.min to /etc/conf/dfile.
NOTE:     Saved /hp-ux in /SYSBCKUP.
        * Generating a new kernel.
        * Following is output from "/etc/config":
          ...
```

If you see these messages you can be sure that the update process copied your kernel. Doing this again will overwrite the backup copy. You may want to overwrite the old kernel if you have updated from 7.0 to 8.0, for example. Usually this step is done prior to kernel reconfiguration, which is covered in a separate module in this course.

Make a Full Backup

This will be covered in a separate module in this course. The step is an important one, since you can always get your file system back to this point in the future should there be a serious problem with your disk or file system.

Make a Recovery System

This will be covered later in this appendix. Making a recovery system applies only to Series 300/400/700 systems. The step is also an important one, since you can use the recovery system in the future to recover from a serious problem with your disk or file system.

B-16. SLIDE: After the Update — Instructor Notes

Purpose

To explain the steps students should take after updating their system.

Key Points

- Check `/tmp/update.log` and the `/etc/newconfig` directory.
- Check to see if you need to install a new versions of the files in `/etc/newconfig`.
- You only need to make a new backup kernel under certain conditions. Be sure to read the update.log file described above.
- Perform a full backup. This will be covered in a separate module in this course.
- On a Series 300, make a new recovery system. You cannot use a recovery system from a previous version of the operating system to recover your system. For example, if you are running HP-UX 8.0, and your recovery system was made using HP-UX 7.0 you cannot use the recovery system to recover from a problem.

References

A new manual at 8.0 that may be helpful when problems arise is: *Solving HP-UX Problems* HP P/N B1862-90010

An earlier version of this book was: *Troubleshooting HP-UX Systems* HP P/N 92453-90026

Transition

The next slide changes the pace a bit to complete the objectives stated earlier that we would cover steps necessary to configure a system as a network server.

B-17. SLIDE: Configuring for a Network Update

Configuring for a Network Update

- Verify correct network operation between client and server.
- Configure one host as an update server.
  - Create /netdist directory.
  - Using updist, load the software.
  - Set up public network access.
  - Put the TOOL fileset in a public directory.
  - Start the network update server daemon, netdistd.

50722 B-17 145

Student Notes

The server should have already been configured for correct LAN operation before configuring as a network update server. Nevertheless, you should verify correct network operation on client and server. The specifics are beyond the scope of this class, but a cookbook version of the LAN setup process is described in an appendix to this course.

Configure one host as an update server. The server system needs a directory created to hold an image of the files from the distribution media. Sufficient disk space is required to hold the duplicate copies of all the files from the update media. This can be a substantial amount of disk space depending upon the software products expected to be supported for network update.

The steps to configure an update server follow.

Create /netdist directory

The server system needs a root level directory /netdist created to hold an image of the files from the distribution media.

Ensure that sufficient disk space is available to hold copies of all the filesets from the update media plus additional space for overhead information about filesets, dependencies, sizes, etc. This can be a substantial amount of disk space depending upon the software products expected to be supported for network update. Check the README notices for sizes of the filesets. The filesets account for the largest share of the disk space.

The updist program will help with an option to check disk space before actually uploading files.

Using updist, load the software

The /etc/updist command will load the filesets from the UPDATE media onto the server's disk. The behavior of the updist program is similar to update, except that the files will be loaded into a special directory, and no customization is done after the process completes. At the completion of the updist program, the file system of the server remains unchanged except for the new (and very large) directory, /netdist.

The Main Menu of updist is very similar to the update main menu with a few minor changes:

All references to update are now to updist
Destination is now the /netdist directory vs the / directory.
All references to the destination will now be to a directory /netdist rather than the server's file system / as in update.

Follow the same steps to select desired filesets and partitions and the like as you would with update. Refer to the summary slide earlier for a review of the steps, mentally substituting updist for update.

Set up public network access

A user login that allows read access to the filesets, but restricted access to the system is needed by client systems that wish to update using the netdist server. Because of the built-in security in the anonymous ftp user login, this account is a good choice for this type of access. The ftp login does not usually require a password be provided, hence the name "anonymous ftp". Setup of "anonymous ftp" can be done with a few key presses using SAM if your system is already up and running on the network.

The specific steps are recounted here. From the "System Administration Manager" main menu, select "Networks/Communications", then select "ARPA Services Configuration", and finally select "Create Public Account for File Transfers". Make sure you allow remote users to retrieve files when prompted by SAM. This will allow access to /users/ftp/dist.

Put the TOOL fileset in a public directory

The TOOL fileset from the update media should be copied into /users/ftp/dist on the server. This allows the remote system to use the ARPA network service ftp to extract the update program from the server before running update.

For a server that intends to serve Series 300 or 400 systems, in the directory: /users/ftp/dist you will put a copy of the TOOL fileset with the command:

```
# cd /netdist/300/TOOL/product
# tar cvf  /users/ftp/dist/TOOL.300  etc system
# chmod 444  /users/ftp/dist/TOOL.300
```

Replace 300 with 800 or 700 as appropriate. A network updist server can support all three platform's software if disk space is sufficient.

Start the network update server daemon, netdistd

This cannot be done in the startup file, so you will have to type in the command manually:

```
# /etc/netdistd -l    # use a lowercase "L"
```

B-17. SLIDE: Configuring for a Network Update Instructor Notes

Teaching Tips

The process is a bit complicated unless the students have a networking background. Try to present the steps outlined in the student notes as more of a COOKBOOK approach.

Avoid getting bogged down in the details.

This will give the general overview of the process of configuring a network update server.

Key points

- Networking must already be running on client and server
- Only one host is configured as an update server, multiple updates can occur at the same time.
- Use `updist` on server to upload files.
- Anonymous `ftp` access is needed on the server.
- The TOOL fileset must be copied to a public directory
- The `netdistd` daemon must be started.

For the 8.0 release, you will probably need as a **minimum** 83 megabytes for a series 300/400 and 75 megabytes for a series 600/800. For a mixed cluster server this space requirement could reach 400 megabytes when both series 300/400 and series 600/800 HP-UX operating systems and additional applications are supported.

Some additional steps that are needed were not covered in the interest of time. These include:

- Check `/etc/services` file for an entry for netdist server There must be a line like:

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

- Check `/usr/adm/inetd.sec` file for access by remote systems on the server.

References

The following manuals are excellent references for this appendix, you should have them on hand when preparing to teach this course and during the class as a reference.

Installing and Updating HP-UX HP 9000 series 300/400 HP P/N B1862-90002 12/90

Installing and Updating HP-UX, HP 9000 Series 600/800 HP P/N B2437-90003 12/90

B-18. SLIDE: Performing a Network Update

Performing a Network Update

- Verify connectivity to server system
- Load the TOOL fileset using ftp to client
- Run update on client system
- Complete update process as before

50722 B-18 146

Student Notes

Updating across the LAN is really just as simple as using local physical media to perform an update. You can only `update` your system using this method, you cannot `install` a system across the network.

The steps required to perform a network based update follow.

Verify connectivity to server system.

Login to the remote system that is the designated netdist server to verify the connectivity, and ensure that the netdist daemon is running.

Load the TOOL fileset using `ftp` to client.

Copy the Tools fileset across the network from the server. The steps from a Series 300/400 client side are:

```
# ftp netdist-server
login:   ftp
passwd:  anonymous
ftp>  get dist/TOOL.300 /tmp/TOOL
```

... file is copied ...

Next, restore the update tools into the proper directories:

```
# cd /
# tar xvf /tmp/TOOL
```

Run `/etc/update` on client system.

Note that this is `update` on the client system, *not* `updist` which is used only in configuring the server system.

In Update Main Menu, select "Change Source or Destination"

In "Change Source or Destination" Menu, select: "From Netdist Server to Local System"

You will have to interact with pop-up windows, and give required information such as codewords, host names of the server, etc.

Return to the Update Main menu and select one of: "Select All Filesets on the Source Media" or "Select Only Filesets Currently on your System" or "Select/View Partitions and Filesets"

Mark the desired filesets and partitions to load.

Press **Start Loading**.

Wait for completion.

Complete update process as before. Proceed to post-update procedures described above for the local media based update process.

B-18. SLIDE: Performing a Network Update Instructor Notes

Purpose

To describe the process of across the network update.

Teaching Tips

The *netdist-server* host name will be a mystery to some students, just explain that it is a unique network identifier for the system. This is not a network class, and many students will not recognize the command. It should not be stressed, just indicate that it is another way to copy files.

In the commands for downloading the TOOL fileset, the reference to

```
ftp> get dist/TOOL.300 /tmp/TOOL
```

for series 300/400 systems, must be changed to

```
ftp> get dist/TOOL.800 /tmp/TOOL
```

for series 600/800 systems.

Lab possibilities

In most class room setups the students will not be able to do an actual update, on the Series 600/800 there is a simulation program that gives the student some idea of how an update goes.

Appendix C — Setting Up a Local Area Network

About This Appendix

This appendix will provide a brief overview of the Local Area Network product. It will provide you information on the following:

- The functionality of the LAN.
- The functionality with respect to various hardware platforms.
- The steps for installing a LAN.

Overview of Module C

This optional appendix is a quick overview of the LAN product. It is not intended to be a complete tutorial on LAN and its capabilities. If interested, there is a separate course on the LAN and ARPA/Berkeley services.

This appendix is identical to the appendix in the 5-day Sys Admin class, 51436 and 51482.

C-1. SLIDE: LAN Functionality

LAN Functionality

- File Transfer
- Remote Login
- Electronic Mail
- Remote Command Execution
- Remote Interprocess Communication
- Remote File Access

50722 C-1 147

Student Notes

The software packages for LAN include:

ARPA	Advanced Research Projects Agency networking package
Berkeley	University of California at Berkeley networking package
NS	HP Network Services
OSI	Open Systems Interconnection
NFS	Network File System

These software packages include the following capabilities:

Summary of Networking Capabilities on Series 800

Capability	ARPA	Berkeley	NS	OSI
File transfer	ftp	rcp	dscopy	FTAM
Remote Login	telnet	rlogin	vt	VTAM
Electronic Mail	smtp	sendmail	--	X.400
Remote Command Execution		remsh		
Interprocess Communication		sockets	NS-IPC	XTI
Remote File Access			--	FTAM

Note

Until release 8.0, the NS product for the HP 9000 systems included Network File Transfer (NFT), Virtual Terminal for HP3000 (VT3K), and Remote File Access (RFA). With release 8.0 of NS, the RFA features of NS is obsoleted. This feature's capabilities can be replaced with those of the NFS product which has been available since HP-UX 6.0.

Another note: ARPA, Berkeley and NFS require LAN/9000 Link or X.25 Link products.

NFS Services

Another important networking package available on HP-UX is NFS. Indeed, NFS is the defacto UNIX remote file access standard today for networked file systems. NFS was developed by Sun Microsystems.

NFS allows a client node to perform transparent file access over the network. By using NFS, a client node operates on files residing on a variety of servers and server architectures, and across a variety of operating systems. File access calls on the client (such as read requests) are converted to NFS protocol requests and sent to the server system over the network. The server receives the request, performs the actual file system operation, and sends a response back to the client.

A client gains access to that file system using the `mount` command to request a file handle for the file system. The client must mount the file system directly from the server on which it resides.

The features of NFS can be best explained by listing the goals of NFS:

- *Transparent access to a remote file system.*

  The goal here is to allow access to files and directories on a remote file system and to access a remote file the same as a local file.

- *Machine and operating system independence.*

  The NFS protocol can be implemented on non-UNIX systems. The tradeoff of this is non-support of operating system specific functionality.

- *Good recovery characteristics.*

  A machine offering a file system is not sensitive to a crash of a machine using its file system.
- *Easy portability.*

  NFS has been ported to several operating systems and hardware architectures.

Examples of features offered by NFS that ARPA/Berkeley and NS networking services do not offer are as follows:

- Mirror login environments of local users on a remote machine (VHE).
- Send message to all users on a remote machine (`rwall`).

Note that the 8.0 release of HP-UX includes two additions to ARPA services which involve remote systems. `bootpd` and `tftp` are new to the ARPA product in 8.0. Users now have the capability to set up a "bootpd" server and add "boopd" clients to be served by the bootpd server. (A bootpd server is a remote system which can boot from another system using the Bootstrap Protocol (BOOTP).) `tftp` is the user interface to the Trivial File Transfer Protocol that allows users to transfer files to and from a remote machine. SAM supports these new capabilities.

Of course, NFS duplicates some of the functionality provided by ARPA/Berkeley and NS. For example, NFS offers its equivalents of the `remsh` command (`rexd`), the `rwho` and `ruptime` commands (`rusers` and `rup`), and the `ping` command (`spray`). Students should note that they can get similar functionality for most of the NFS commands that duplicate the ARPA/Berkeley/NS services for less overhead. For this reason, we recommend that students use the ARPA/Berkeley commands as listed in the earlier table when possible.

C-1. SLIDE: LAN Functionality — Instructor Notes

Key Points

Basically, the software products described above provide a number of different ways to do the same types of tasks. Simple file transfer and remote login capabilities are present in all forms. The difference is in what vendor's hardware supports which service.

OSI is shipped by HP in a software package called OTS, for OSI Transport Services. FTAM stands for File Transfer, Access, and Management. XTI stands for X/Open Transport Interface.

New to Release 8.0

- New ARPA services for remote machines

  The Student Notes contain brief descriptions of two new commands added to ARPA functionality: `bootpd` and `tftp`.

- New implementation of BSD/ARPA networking

  Release 8.0 of HP-UX contains a new implementation of BSD/ARPA networking, based on 4.3bsd. Earlier releases of HP-UX (prior to 8.0) were based on 4.2bsd networking. 4.3bsd has become the defacto standard for networking on UNIX systems. This change has introduced a few small incompatibilities at both the object and source code levels with earlier releases of HP-UX. Most of these changes will impact programmers, not system administrators.

  Here are a few changes/incompatibilities which could affect system administrators. The syntax for the `ifconfig` and `netstat` commands have changed, and a new command, `lanconfig`, has been added. HP's current versions of these commands are based on 4.2bsd, and they are very similar to the 4.3bsd commands. These changes may not be noticed by most customers, but customers who use these commands in scripts may need to modify those files.

  - `ifconfig` command

    The HP-UX 7.0 command allowed for configuring the link level encapsulation (ieee, ether) on the `ifconfig` command. In 8.0, the new `lanconfig` command should be used for this purpose.

  - `netstat` command

    New options have been added to the `netstat` command. The `-I interface` option can be used to show information for a specified interface. The `-f address_family` option can be used to show the statistics and address control block reports for only the specified family.

  - `lanconfig` command

    The `lanconfig` command has been added to configure network interface options, for example, ieee and ether encapsulation. Refer to the man page `lanconfig(1m)` for details.

- Removal of RFA from NS Services

  All commands which were impacted by RFA have been modified to function correctly without this feature (for example, shells, `mount`, backup utilities). In SAM, for instance, the configuration of RFA is removed from the SAM menu tree.

- Network Tracing and Logging

  Prior to 8.0, most of the networking products used their own tracing and logging mechanisms. At 8.0, a common utility has been adapted by a wide range of networking products, including the links (802.3,X.25), MAP, GOSIP, ARPA Services, and so forth. This common tracing will run on all HP 9000 series systems.

C-2. SLIDE: A Sample LAN Configuration

A Sample LAN Configuration

What the network looks like:

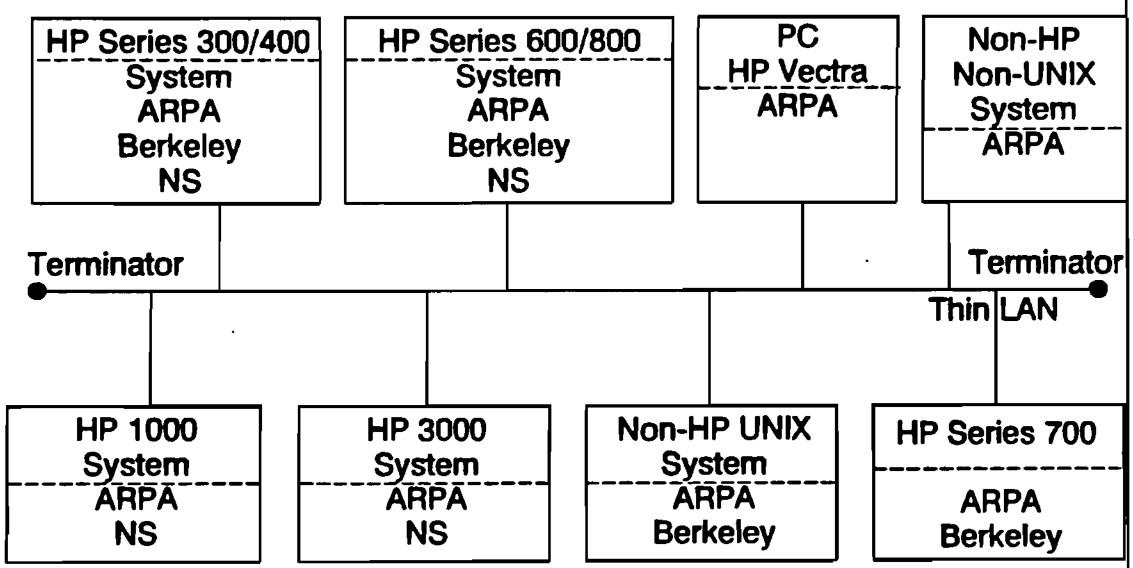

50722 C-2 148

Student Notes

Note

This is a sample of what your LAN network may look like. The slide uses generic systems for the non-HP entities. You need to refer to the appropriate vendors specifications for supportability of these LAN products.

There ARE exceptions, but in general the following applies:

ARPA	Berkeley	NS	OSI
HP or non-HP	HP or non-HP	HP plus VAX	HP or non-HP
UNIX or non-UNIX	UNIX	UNIX or non- UNIX	Unix or non-UNIX

This means that generally, NS is limited by proprietary nature. As such, it runs on HP equipment only (with a few exceptions). Berkeley is generally limited to UNIX-like operating systems. ARPA is, by far, the most widespread utility, spanning product vendors as well as operating systems. OSI is still gaining acceptance, but will play an increasingly important role as an industry standard.

C-2. SLIDE: A Sample LAN Configuration

Instructor Notes

Purpose

To present a sample LAN configuration, implementing three or more of the available LAN services.

C-3. SLIDE: Manual Steps for Installing a LAN

Manual Steps for Installing a LAN

1. Connect the LAN cable to the LAN connector.
2. Ensure that the LAN driver is in the kernel (if not, add it).
3. Assign host name in `/etc/rc`.
4. Verify LAN Internet address in `/etc/netlinkrc` and /etc/hosts.
5. Edit `/etc/hosts` file and add a line for each system on LAN.
6. Reboot and start up the network.
7. Verify the configuration.

50722 C-3 149

Student Notes

More information on each step is provided below.

Before you start this process, you need to gather some information:

- Logical name for the lan card (lan card ID).

  The logical name of the LAN card is determined differently on the various HP 9000 systems, as follows:

  - On HP 9000 Series 300s, 400s, and 700s, LAN card lan0 is the LAN card with the lowest select code. LAN card lan1 is the LAN card with the next higher select code, and so forth.
  - On HP9000 Series 800s and 600s, the LAN card's numbers are determined by their slot numbers instead of select codes.

- Hostname for the system

  This is the name by which people using the network can uniquely identify the system you are adding. You can obtain the system's official hostname from the person who administers that system or the network containing that system. For example: hpscad.

- IP address for the lan card,

  The IP address allows the LAN card being added to uniquely identify messages sent from it and recognize messages sent to it. You can obtain the IP address from the person who plans your network's Internet addressing scheme. This person derives a set of unique IP addresses for the systems on your LAN from a unique network address obtained from HP or a higher Internet authority. An example of an IP address is: 192.6.12.33

- Connect the LAN cable to the LAN connector. Terminate both ends with terminators. (We assume that your LAN card is already installed in the system. If not, you must install the card.)

- Insure that the LAN driver is in the kernel (if not, add it).

  If you are on a Series 300/400, check for the following lines in your `dfile`:


```
lla
lan01
nsdiag0
```

  If these lines are not in your `dfile` add them in. You also have the option of copying the file `/etc/conf/dfile.full.lan` to your current `dfile`. Once you add these lines, you will need to generate a new kernel. A review of this is shown below:

```
# cd /etc/conf
# cp dfile.full.lan dfile
# make -f config.mk
# mv /hp-ux /SYSBCKUP
# mv hp-ux /hp-ux
```

  If you are on a Series 600/800, check for the following lines in your S800 input file:

```
include netdiag0;
include lan0;
include lan;
```

  If these lines are not in your S800 file add them in.

Note

The `lan0` and `lan1` drivers are among those which are automatically configured by `hp-ux`, so you do not need to specify their hardware and driver paths (io statement) in the S800 file.

Once you add these lines, you will need to generate a new kernel. A review of this is shown below:

```
# cd /etc/conf/gen
# cp S800 S800.BCKUP
# uxgen S800
# cp /hp-ux /SYSBCKUP
# mv ../S800/hp-ux /hp-ux
```

- Assign host name in `/etc/rc`. Look for the line:

  ```
  SYSTEM_NAME=unknown
  ```

  Change unknown to your *hostname*. You could also assign host name as follows:

  ```
  hostname 'uname -n'
  ```

  The output of the `uname -n` command is the name of your system.

- Configure LAN Internet address in `/etc/netlinkrc`. The `netlinkrc` file is commented extensively and explains where to change or add information. You must change the `ifconfig` command to match your internet address. For example:

  ```
  /etc/ifconfig lan0 inet 15.3.43.234 up
  ```

  If the line `/etc/ifconfig lan0 inet 'hostname' up` exists in the `/etc/netlinkrc` file, you don't need to change anything. (This should be the case at HP-UX 8.0.) If you are subdividing networks into sub-networks, you must use the netmask option on the ifconfig command line.

  You must also change the `nodename` command to define your system NS nodename. Look for the following command:

  `/bin/nodename` *nodename.domainname.orgname*

  Change *nodename*, *domainname*, and *orgname* to match the name of your node, domain, and organization.

- Edit `/etc/hosts` file and add a line for each system on LAN. The `/etc/hosts` file is the host name database. It associates internet addresses with official host names and aliases. It allows a user to refer to a host by a symbolic name instead of an internet address. An example `hosts` file is shown below. You would add your internet address and nodename to this file.

  ```
  15.13.167.84   hpserv    #HP-UX 9000/350
  14.3.33.100    hptest    #HP-UX 9000/370
  15.3.43.234    hpclone   #HP-UX 9000/825
  ```

- Reboot and start up the network. Use `shutdown -r` (or `reboot` if you are in single-user mode) to reboot your system. If you generated a new kernel in step 2 above, the changes will now take effect. Also, by rebooting you cause the network to be set up.

- Verify the lan is installed and powered up.

After you have rebooted, it is wise to verify that the lan is indeed installed and powered up. There are several commands which can be used on HP 9000 systems to retrieve this information:

Commands Used to Verify LAN Setup

Command	Information Gained ...
`ping`	Percentage of packets transmitted and received between systems; lists IP address of system being accessed.
`lanscan`	Status of lan (up/down), encapsulation methods, select code/hardware path, lu number, net name.
`ifconfig`	Status of lan (up/down), IP addresses of two systems, net name
`verify`	600/800 only; lu and h/w path numbers, driver name, hardware and software status
`telnet`	Access a system using ARPA service; if successful, will receive `login` prompt
`rlogin`	Access a system using Berkeley service; if successful, will receive `login` prompt

Descriptions and examples of using each of these commands are presented below.

— `ping`

`ping` sends ICMP ECHO_REQUEST packets to network hosts. Each packet that is echoed back represents a successful transfer of information between the two systems. The command `ping hpasbb` issued on a Model 835 to test the network connection between it and a Model 345 with the host name of `hpasbb` and the IP address of 15.2.92.102 produced this output:

```
64 bytes from 15.2.92.102; icomp_seq=0. time=5. ms
64 bytes from 15.2.92.102; icomp_seq=2. time=6. ms
64 bytes from 15.2.92.102; icomp_seq=3. time=5. ms
                        .
                        .
                        .
64 bytes from 15.2.92.102; icomp_seq=28. time=4. ms
64 bytes from 15.2.92.102; icomp_seq=29. time=5. ms
64 bytes from 15.2.92.102; icomp_seq=30. time=5. ms

-------hpasbb PING statistics ------
31 packets transmitted. 31 packets received. 0% packet loss
round-trip (ms) min/avg/max = 4/5/6
```

— `lanscan`

`lanscan` displays LAN device configuration and status. The command `lanscan` issued on a Model 345 produced this output:

```
Select    Station          Dev  Hardware    Net-Interface  Net-Mgt  Encapsulation
Code      Address          lu   State       NameUnit State ID       Methods
21        0x08000906F908   0    UP          lan0     UP    2        ETHER IEEE8023
```

– ifconfig interface

Where interface is the LAN card id. For example, the command ifconfig lan0 issued on the same Model 345 as above produced this output:

```
lan0:  flags=63<UP,BROADCAST,NOTRAILERS,RUNNING>
         inet 15.3.43.200 netmask fffff800 broadcast 15.3.47.255
```

– verify

There is an additional system verification command which can be used on Series 600s/800s only. You can examine the I/O configuration of your system by typing the command ioscan -f. You will get a listing on the screen that resembles the following example:

```
Class      LU  H/W Path  Driver                        H/W Status       S/W
Status
=========================================================================
processor  -    0        processor                     ok(0x0)          ok
cio        -    4        cio_ca0                       ok(0x1000)       ok
                                   .
                                   .
                                   .
lan        1    4.4      cio_ca0.lan0                  ok(0x6)          ok
memory     -    12       16Meg_memory                  ok(0x800)        ok
```

If you had added a LAN card in slot 4 of the channel adaptor in a Model 835, you would see that lan card included in the system configuration. The line that reads,

```
lan   1  4.4   cio_ca0.lan0    ok(0x6)     ok
```

shows that.

– Using telnet and rlogin

Of course, the final test for your lan installation is the ability to access other systems. Depending on which type of networking service you are using, you should now try to test the link between the system configured and other systems listed in its /etc/host table.

For example, your /etc/host table contains the following entries:

```
15.13.006.3   mysystem     # describes system who owns this host table
15.13.201.8   edsbox       # describes Ed's system
15.13.011.7   marysbox     # describes Mary's system
```

If you are using ARPA Services, you could access Ed's system by:

```
telnet 15.13.201.8
```

If you are using Berkeley Services, you could access Ed's System by:

```
rlogin edsbox
```

In both cases, a successful access results in receiving the `login` prompt on your display.

Note

Be aware that the various services have security layers which may have been built in to the system you are trying to access. This may prevent you from successfully accessing a system for which you have correctly set up access. For example, if Ed is using ARPA Services, he may have set up on his system a file named `/usr/adm/inetd.sec`. If he has specified a line in this file that, for example, disallows access from any system that has a telnet access starting with the numbers "15.13", you will be denied access to his system.

You will want to explore the security features of the various networking services to find one that fits your environment's needs.

C-3. SLIDE: Manual Steps for Installing a LAN — Instructor Notes

Purpose

Walk through the process of manually adding a lan subsystem to an HP 9000 system.

Key Points

Be sure to point out to students the possibility of correctly installing and configuring a lan subsystem and card(s) on a system, and not being able to access other systems because of security measures implemented on the system being accessed.

C-4. SLIDE: Network Configuration with SAM

Network Configuration with SAM

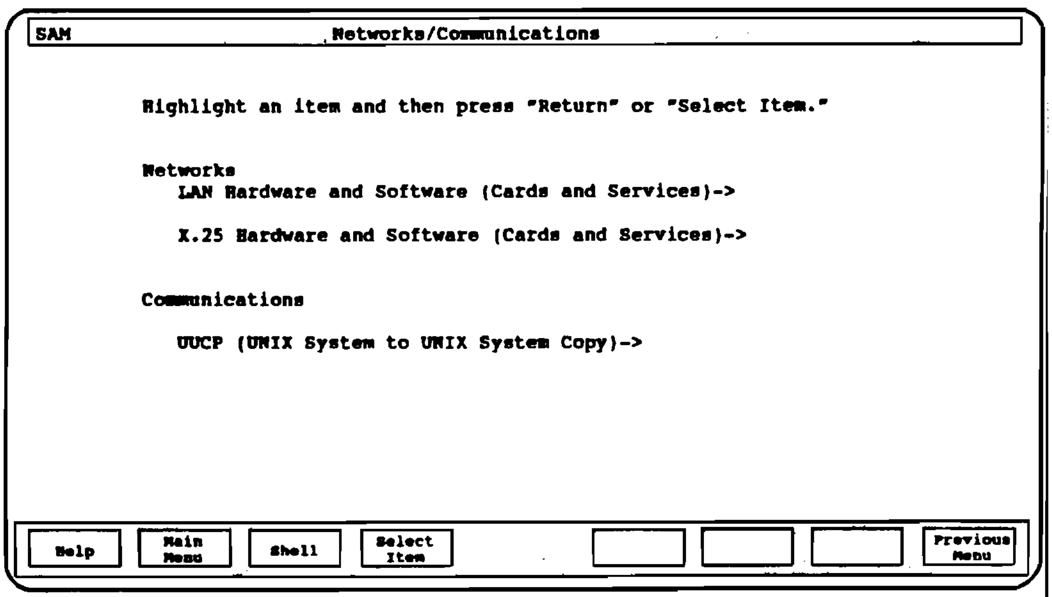

50722 C-4 150

Student Notes

You can use SAM to perform many network administration tasks. When you choose "Networks/Communications" from SAM's Main Menu, you get this menu.

Under "LAN Hardware and Software(cards and Services)," SAM will:

- Add a new LAN card
- View/Modify a LAN Card's Configuration
- Power Up an Existing LAN Card
- Power Down an Existing LAN Card
- Provide ARPA Services Configuration
- Provide NFS (Network File System) Configuration

- Provide NS (Network Services) Configuration
- View/Modify This System's NS Node Name

Under "X.25 Hardware and Software(cards and services)", SAM will:

- Add a New X.25 Card
- View/Modify a X.25 Card's Configuration
- Delete an X.25 Card
- Power UP an Existing X.25 Card
- Power Down an Existing X.25 Card
- Provide ARPA Services Configuration
- Provide NS (Network Services) Configuration
- Provide PAD Services Configuration
- View/Modify This System's NS Node Name

Under "UUCP (UNIX System to UNIX System Copy)", SAM will:

- Add a Device
- Remove a Device
- View/Modify Device Configuration
- Add a System
- Remove a System
- View/Modify System Configuration
- Modify Permissions when Remote Systems Call in
- Modify Permissions when Local Systems Call out

C-4. SLIDE: Network Configuration with SAM Instructor Notes

C-5. TEXT PAGE: References

Reference Listing

For more complete information on installing and configuring LAN for HP-UX 8.0, see the following:

- *Finding Information, HP 9000 300/400*, P/N B1862-90001
- *System Administration Tasks, HP 9000 Series 300*,P/N B1862-90008
- *How HP-UX Works,Concepts for the System Administrator, HP 9000 Series 300/400 & 600/800*,P/N B1862-90006
- *How Remote Access: User&'s Guide, HP 9000 Series 300 & 800*, P/N B1862-90011
- *Using Network Services, HP 9000 Series 300/400 & 600/800*,P/N B1012-90009
- *Inst. & Admin. Network Services, HP 9000 Series 300/400 & 600/800*,P/N B1012-90010
- *Networking Overview, HP 9000 Series 300/400 & 600/800*, P/N B1012-90012
- *Using NFS Services HP 9000 Series 300/400 & 600/800*, P/N B1013-90008
- *Inst. & Admin. NFS Services, HP 9000 Series 300/400 & 600/800*, P/N B1013-90009
- *Programming & Protocols for NFS, HP 9000 Series 300/400 & 600/800*, P/N B1013-90010
- *Using ARPA Services, HP 9000 Series 300/400 & 600/800*, P/N B1014-90006
- *Inst. & Admin. ARPA Services, HP 9000 Series 300/400 & 600/800*, P/N B1014-90007
- *HP OSI Express 802.4 Hardware Inst. & Config. Guide, HP 9000 Series 300/400 & 600/800*,P/N 28604-90001
- *HP FTAM/9000 Ref. Manual, HP 9000 Series 300/400 & 600/800*, P/N B1033-60500
- *HP FTAM/9000 Progmr. Guide, HP 9000 Series 300/400 & 600/800*, P/N B1033-60510
- *HP FTAM/9000 User&'s Guide, HP 9000 Series 300/400 & 600/800*, P/N B1033-60520
- *OSI Config. & Trouble Guide, HP 9000 Series 300/400 & 600/800*, P/N 32069-60001
- *Inst. & Admin. OSI Transport Services, HP 9000 Series 300/400 & 600/800*,P/N 32069-60003
- *XTI Programmer&'s Guide, HP 9000 Series 300/400 & 600/800*, P/N 32069-60002

C-5. TEXT PAGE: References

Instructor Notes

Note

The part numbers in the reference listing may be out of date.

Solutions

2-20. LAB: Exercises

1. List all of the SCSI devices connected to your lab system (if any). Include the bus addresses of the devices in your list.

Answer:

Systems will vary but typical SCSI devices include disk drives, DDS/DAT tape drives, CD-ROM drives, and magneto-optical disk drives.

Addresses of the devices can be found on the back of the device.

Series 800 students can try the `ioscan -u` command to list devices currently connected to the system.

2. List all of the HP-IB devices connected to your lab system (if any). Include bus addresses in your list.

Answer:

Typical HP-IB devices include disk drives, DDS/DAT tape drives, cartridge tape drives, CD-ROM drives and possibly a printer.

Addresses of the devices can be found on the back of the device.

3. What is the bus address of your system disk (the one the system booted HP-UX from)? What is its device file name?

Answer:

If the system disk is SCSI its bus address should be 6.

If the system disk is HP-IB its bus address should be 0.

You can run the `bdf` command to see which device file represents the "/" directory. This can be verified by executing `ll` on the device file and looking at the minor number.

4. List any other devices that are attached to your system. Make sure you inlcude the interface type of each device and port number if needed.

Answer:

Answers will vary based on type of system and Education Center peripherals. Ask your instructor to verify your list.

5. Draw a diagram showing the I/O interfaces and peripheral devices on the system that you administer or use at your own work site.

List all of the device files for the peripherals on your system.

Answer:

Answers will vary.

6. Series 800 only:

Invoke the `ioscan -uf` command and look at the output. What are the hardware addresses of the system disk, tape drive (if any) and MUXs?

Answer:

Output of `ioscan -uf` will show the hardware addresses of all devices recognized by the system you are on (this includes adapters, modules, I/O interfaces and external peripheral devices). This command is useful if you want to get an overall picture of the hardware configuration of your system.

Ask the instructor for the specific hardware addresses for the system being used in the classroom.

7. Series 800 only:

Run the following command and describe the output:

```
# lssf /dev/dsk/*
```

Answer:

The output is the complete description of all device files in the `/dev/dsk` directory. This command can be very useful for locating device files on the Series 800.

3-14. Worksession: Review Questions

1. What is the difference between executing `shutdown` and executing `init s`?

Answer:

They both take you to single user mode but `shutdown` also takes the time to warn users, wait a specified amount of time, kill all unnecessary processes and unmount the file systems. `init s` is a single-user mode but it is not the quite single-user mode you get with `shutdown`. In fact `init s` is just one part of the `shutdown` procedure.

2. Why might users be angry with you if you used `reboot` to reboot the system instead of `shutdown -r`?

Answer:

`reboot` does not warn users of the shutdown or give them any time to properly log off. If users had applications running and work not saved, `rebooting` the system could cause them to loose their work.

`shutdown -r`, of course, will give the users a warning and a reasonable time period for the users to properly save their work before it reboots the system.

3. What are some of the things `/etc/rc` does at boot time?

Answer:

`/etc/rc`:

- mounts all defined file systems
- activates defined swap devices
- starts the printer spooler
- etc.

See *The* `/etc/rc` *Program* in this module for complete details.

4. How does `init` know what run-level to enter after booting?

Answer:

From the `initdefault` entry in `/etc/inittab`

3-15. LAB: Exercises

1. What is the default run-level of your system?

Answer:

Look at the `initdefault` entry of the `/etc/inittab` file.

2. What is the current run-level of your system?

Answer:

Use the `who -r` command to find your current run-level.

3. Determine your terminal port then find your terminal's entry in the `inittab` file.

Answer:

`who am i` or `tty` will allow you to see your terminal's name (device file). Scan `/etc/inittab` to find the corresponding `getty` entry.

Note

If you are connected to your system via a DTC (Data Terminal Controller), you will not find an entry in `/etc/inittab`. DTC connections are handled by the `telnet` network service.

4. If you have a graphics display as your system console, put your system in run-level 3 to start HP-VUE.

Answer:

`init 3; exit` works best.

5. How can you get your system from run-level 3 back to the default run-level?

Answer:

`reboot` or `init` *default_run-level* will work.

6. List all of the processes that are running on your system. Can you find your processes in the listing?

Answer:

7. List just the processes you are running from your terminal. Now start a background `sleep` process that we can play with. List your processes again and make a note of the `sleep` PID.

Answer:

```
# ps -f
# sleep 2000 &
# ps -f
     UID   PID  PPID  C    STIME TTY      TIME COMMAND
    root 26013  9471 12 15:50:15 console  0:00 sleep 2000
    root  9471     1  6   Jun 23 console  0:05 /bin/ksh
    root 26014  9471 22 15:50:16 console  0:00 ps -f
```

The sleep PID is 26013.

8. Kill the sleep process then list your processes to ensure it has died.

Answer:

```
# kill 26013
```

9. Shutdown your system to the quite single-user mode using a grace period of 30 seconds.

Answer:

shutdown 30

10. Once in single-user mode list all running processes. How does this compare to the listing you saw in multi-user mode?

Answer:

The listing will only have about 10 processes listed. This is why we call it the **quiet** single-user mode.

11. Reboot your system so it will come back up into the default run-level.

Answer:

reboot

5-12. WORKSESSION: Review Questions

1. Describe the user account information needed to build a user account on the system.

Answer:

See topic "Adding a New User with SAM" for details.

This information is stored in `/etc/passwd` in the following format:
`user_name:password:user_id:group_id:comments:home_dir:startup_program`

2. What is the difference between removing a user and deactivating a user on the system?

Answer:

Reference: "Removing User Accounts" and "Deactivating and Reactivating a User"

Removing a user actually takes the account off the system permanently by removing the user's entry from the `/etc/passwd` file. The only way to get the user back on the system is to add the account again.

Deactivating a user simply changes the user's password to *. The user's account still exists but s/he can not log in.

Both removing and deactivating a user have several actions that may be taken regarding the user's files.

3. List three ways a user's password may be changed.

Answer:

Reference: "Setting a User Password" under "Changing the User's Password"

1. Using SAM's "Change a User Password" option.
2. By the superuser using the `passwd user_name` command.
3. By the user using the `passwd` command.

4. Why is it important to protect directories (especially user home directories) with 755 permissions?

Answer:

Reference: "User Security Considerations"

755 permissions look like this on a directory:

```
drwxr-xr-x  2 ralph    users       1024 Jun 17 15:00 /users/ralph
```

Notice that anyone can list the directory contents (the r permissions) and anyone can cd to, or otherwise use, the directory (the x permissions) but only ralph (the owner) can change its contents (the w permission).

This prevents everyone else from:

removing files from the directory (`rm` looks at directory permissions)
creating files in the directory (like network security files)
renaming files in the directory

5. What is the difference between `/etc/profile` and `.profile`?

Answer:

Reference: "Customizing the User Environment"

`/etc/profile` is the system login file used for all users. Anything that the system administrator puts in this file will have an effect on all users (running `ksh`, `sh` or a restricted shell) when they log in.

The `.profile` file is the user's own file. Any command that s/he adds to this file will be executed when s/he logs in and only when this user logs in, no one else. This is how a user can make his/her environment different from the defaults set in `/etc/profile`.

6. If you were planning to shutdown the system at 8:00pm, how might you communicate that to all of your users?

Answer:

Reference: "Communicating with System Users"

Using the /etc/motd file. Anything that you put into the /etc/motd file will be displayed whenever a user logs into the system.

7. When 8:00pm arrives, how can you immediately warn anyone still logged in that the system will be shutdown now?

Answer:

Reference: "Communicating with System Users"

Using the `wall` command. This will send a message to everyone that is logged into the system.

5-13. LAB: Exercises

1. When it is your turn, invoke `sam` to add a new user account to your system. Remember that you must be logged in as `root` (the superuser) to invoke SAM.

Give the user the following attributes:

- Use your (or your partner's) name as a user name.
- Use the default group and home directory.
- Use the default (Bourne) shell as the startup program (`/bin/sh`).
- Use any comments you would like.
- Don't modify the user defaults.
- Make the user assign their own password when they first log in.

Answer:

```
# sam (Return)

    Users ->

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

Fill in the information a requested on the data entry screen and press the `Perform Task` softkey.

When SAM prompts you for a password, enter `,..` to force the user to set a password the first time they log in.

2. Now, exit SAM and look at the `/etc/passwd` and `/etc/group` files. Do you see the user you added?

Answer:

To exit SAM use the `Main Menu` softkey then `Exit`.

To see the contents of the files use the `more` command:

```
# more /etc/passwd (Return)
```
Contents of /etc/passwd

```
# more /etc/group (Return)
```
Contents of /etc/group

In /etc/passwd you should see a line that starts with the new user's name and contains all of the information you used on the data entry screen. Notice the ",.." that you entered for a password in the second field. If you had assigned a real password to this user it would show up here in an encrypted form.

In /etc/group you will find a line that starts with the group name users and your new user will be listed as a group member.

3. Log out as root and log back in using your new account. Try changing the password on your account to your first name. Does this work ? Why or why not? Set a valid password on the account so you can complete the login.

Answer:

Since we set up the user to set a password at first login, you will enter the new user's name at the login: prompt and immediately be prompted for a new password.

You will not be able to change it to your name because of the password naming restrictions insist on at least one non-alphabetic character.

4. Once logged in, execute ll -a to view the files that were created for the user.

Notice the permissions on the files and directories.

Answer:

The files and permissions should look something like this:

```
$ ll -a
total 12
drwxr-xr-x   2 ralph    users       1024 Jun 17  10:33 .
drwxr-xr-x  55 root     sys         1024 Jun 17  10:33 ..
-rw-r--r--   1 ralph    users        818 Jun 17  10:33 .cshrc
-rw-r--r--   1 ralph    users        347 Jun 17  10:33 .exrc
-rw-r--r--   1 ralph    users        377 Jun 17  10:33 .login
-rw-r--r--   1 ralph    users        382 Jun 17  10:33 .profile
$
```

Note that only the user is allowed to modify the home directory (.) and the login files.

5. While logged in as your user, execute the env command and note some of the important things in the environment line PATH, TERM, LOGNAME, etc.

Answer:

The env command displays your entire working environment. A default user environment will be similar to this:

```
$ env
EDITOR=vi
HOME=/users/ralph
LOGNAME=ralph
MANPATH=/usr/man:/usr/contrib/man:/usr/local/man
PATH=/bin:/usr/bin:/usr/contrib/bin:/usr/local/bin:.
SHELL=/bin/sh
TERM=hp
TZ=EST5EDT
$
```

6. This is an exercise to demonstrate the importance of having the user environment set correctly.

Execute sh to create a new shell we can experiment with. Now execute unset TERM this removes that value that describes your terminal type. Commands like vi will not work correctly without a valid value in TERM. Execute vi afile. What happens? Why?

In vi type :q! to exit back to the shell then type Ctrl d to exit the extra shell we started.

Answer:

```
# sh Return
#
This is the new shell prompt
# vi afile Return
I don't know what kind of terminal you are on - all I have is 'unknown'.
[Using open mode]
"afile" [New file]
:q!
#
# exit Return
#
Now we are back to the original shell
```

As vi doesn't know anything about the terminal capabilities (like number of lines or characters per line) without the TERM environment variable, it assumes you are using a line oriented teletype. Though the appearance in vi's open mode is unfamiliar, we still have the complete vi command set at our disposal. The best way out is :q! and a checking and re-setting of the TERM variable.

7. Log out as your user and log back in as root. Using SAM, deactivate your user.

What happens if you try to log in as the user now?

Log back in as `root` and reactivate the user with SAM.

Answer:

```
# sam (Return)

     Users ->

          Deactivate a User Account ...
```

When you try to log in as the deactivated user, you will always get the `login incorrect` message. This is because the password in `/etc/passwd` has been changed to *. This is a special password that will never match anything the user types.

To reactivate the user select `Reactivate a User Account ...` .

8. Create a second user account using whatever user information you want.

When the user has been added use SAM's `Shell` softkey to start a shell. Now look in `/etc/passwd` and the user's home directory to convince yourself that the user really exists. Exit the shell by typing `exit`.

Now remove the user and all of his/her files from the system. Once the task completes read the message displayed by SAM.

You can later check `/etc/passwd` and the user's home directory and notice that the account is gone.

Answer:

Adding a user is identical to the earlier exercise.

To remove the user:

```
# sam (Return)

     Users ->

          Remove a User Account from the System ...
```

Fill in the user name and specify that you want the user's files removed also.

SAM will complete the task and notify you that the user account is gone but the disk search to remove all of the user's files may take some time. This means the entry from `/etc/passwd` is taken out but it may be a few minutes before all of the user's files are removed.

9. Advanced:

Use SAM to change the shell your user use at login to the Korn shell (`/bin/ksh`).

Answer:

```
# sam (Return)

    Users ->

  View/Modify a User's Account Information ...
```

On the data entry screen, move down to the "Start-up Program" field and type /bin/ksh or use the Help softkey to list the shells so you can select one. When you press Perform Task the shell will be changed.

To test this you can log in as the user and execute env. The SHELL variable will be set to /bin/ksh. You could also execute the ps (process status) command and you would see that you are running ksh.

6-8. WORKSESSION: Review Questions

1. What functions does the spooling system provide, and why are they required?

Answer:

The spooling system manages print requests. They are required for optimizing the usage of limited system resources.

2. How would you cancel your own print request?

Answer:

```
# cancel request_id_number
```

3. How would you cancel a print request owned by someone else?

Answer:

You would have to find the request-id number then execute cancel as normal: cancel *request_id_number*

4. How would the owner of a canceled request know that you had canceled it?

Answer:

By a mail message, automatically generated by cancel.

5. How can you cancel all requests for a printer using one command line?

Answer:

`cancel` *printer_name* `-e`

6. What is the model/interface program used for?

Answer:

This is what sets up the printer to print your requests according to the `lp` command line options. There is a model/interface program for each type of supported printer in the `/usr/spool/lp/model` directory.

7. If you have shut down the spooling system using SAM do any printers continue to print?

Answer:

No. There may be information in the printer's buffer that will continue to print but no more information is sent to the printer from the spooler.

An important side effect is that any jobs that were printing when the spooler was stopped will start again *FROM THE BEGINNING* when the spooling system is restarted.

8. If you have stopped the spooling system using SAM can you still submit requests using `lp`?

Answer:

Yes. Requests can still go into the queue but they will not print until the spooling system is restarted.

9. How can you tell other users that a printer is "broken"?

Answer:

By `disable -r"Printer is broken"` *printer*. You could also disable the printer using SAM but you could not supply the specific reason the printer has been disabled.

10. How can you move print requests from the "printer1" queue to the "printer2" queue?

Answer:

`lpmove printer1 printer2`

11. How can you move one print request (request-id "printer2-10") from "printer2" back to "printer1"?

Answer:

`lpmove printer2-10 printer1`

12. What happens to requests in the queue of a printer that is being removed?

Answer:

The requests are all canceled.

6-9. LAB: Exercises

1. Physically connect a printer to your system (if it is not already connected.)

Answer:

Connect the printer. Make sure to note the interface type and check if the printer is properly configured for that type of interface.

2. Add the printer to your system using SAM.

Answer:

See "Adding a Printer With SAM" in this chapter for details.

3. Exit SAM. Submit a request to your new printer using the `lp` command. Did it print? If not, see if you can figure out why.

Answer:

```
lp /etc/passwd
```

4. Check the status of your printer. Is your request listed?

Answer:

```
# lpstat -t
```

Your request will probably not be listed because it already printed.

5. Disable the printer. Now send some output to the printer. Now can you see your queued request?

Answer:

```
# disable -r"Printer down for testing" your_printers_name
# lpstat -t
```

The printer will be marked as "disabled" and your request should now be listed in the queue.

6. Submit more requests to the disabled printer. Now cancel one of the requests from the queue. Check that the requests are gone.

8-19. WORKSESSION: Review Questions

1. What are three benefits of disk sectioning on a Series 800?

Answer:

1. smaller, more manageable file systems on large disks
2. better flexibility in file system layout
3. better control over file system usage

2. What are the required areas on a bootable disk on a Series 300/400? Series 700? Series 800?

Answer:

Series 300/400: 8Kb boot area, root file system, primary swap space.

Series 700: 8Kb LIF directory, root file system, primary swap space, 2Mb boot area.

Series 800: 2Mb boot area, root file system, primary swap space.

3. What is the difference between the "*Add a Hard Disk*" screen and the "*Add a local file system*" screen in SAM?

Answer:

"*Add a Hard Disk*" allows you to define a whole disk with either or both file system and swap space. "*Add a Local File System*" is only used to add an existing file system to the configuration. (On the Series 800 "*Add a Local File System*" will also let you rebuild a file system on the section(s)).

4. What is the difference between the `View File System Information` option and the `View Disk Space Information` option on the "*File Systems*" screen?

Answer:

`View File System Information` lists the mount point directory of the file system where the `View Disk Space Information` option lists the hardware address of the device and any file system swap information.

8-20. LAB: Exercises

1. If there is another disk available for your use, add that disk to your system configuration using SAM. Make sure you create a new file system on the device. Specify any mount point directory you would like. Make sure the disk is defined for automatic mounting at boot time. If you are on a Series 800, ask the instructor if there is a particular disk section you should use.

Answer:

See the topic *Adding a Hard Disk* for your type of system for the complete details. The procedure will vary based on the system type used in class.

2. Use SAM to display the new file system information.

Answer:

```
# sam
    File Systems ->
  View File System Information
```

You should notice the device name and mount point of your new file system and all of its usage information. Also note that the "Mount Now" field is marked with a "y" meaning it is currently available for use. The "Mount on boot" entry should also be marked "y".

3. Unmount the file system you created and view the file system information again. What has changed?

Answer:

```
# sam
    File Systems ->
  Modify a Local File System ...
```

Mark the "Mount Now" field with a n and press **Perform Task** to unmount the file system.

If you `View File System Information` again, you will notice that the "Mount Now" field has changed to "n".

4. Mount the new file system again and add some file system swap to it. Specify a minimum amount of 0 and whatever maximum amount you like.

Answer:

```
# sam
    File Systems ->
  Modify a Local File System ...
```

Answer:

```
# cp your_name fileX
# ls -i your_name fileX
```

The two files will not have the same inode numbers because they are two distinct copies.

5. Hard link another file to `fileX` and call it `fileY`. Use `ll -i` to see the inode number of `fileY`. What happens to `fileY` if you remove `fileX`?

Answer:

```
# ln fileX fileY
# ls -i file?
```

The inode number is the same for all files that are linked together. Notice also that the link count is 2.

Removing `fileX` will have no effect on `fileY` except that the link count will decrease to one.

6. On which file systems would the command `fsck -p` perform a file system check if you did not specify a file system on the command line?

Answer:

The command `fsck -p` would search the `/etc/checklist` file for the file systems to be checked. The first field is the device file name of the file system and the sixth field indicates the *pass number* or the order in which to check the file systems. The root file system should have a pass number of 1. All other file systems should have a higher pass number. The significance of this is that file systems on different disk drives can be checked simultaneously if their pass numbers are equivalent in the `/etc/checklist` file. This can save quite a bit of time depending on your configuration.

7. If applicable, run `fsck` on the secondary file system you created and mounted in the module on creating file systems. Use the preening mode of fsck. Note any problems you may encounter and how to fix them.

Answer:

```
# umount /dev/dsk/your_device_file_name
# fsck -p /dev/dsk/your_device_file_name
```

where *your_device_file_name* might be something like `5s0` on a Series 300/400/700 and `c1d0s5` on a Series 800. Of course these names depend on the bus address and section number of the file system you are using.

The description of any errors can be found in the *Error Message Catalog*.

Note that if you used `-P` the file system probably would not be checked because it was properly unmounted.

8. Properly shutdown your system and run fsck on your root file system disk. Again, note any problems you may encounter.

Answer:

```
# shutdown 0
.
.                 (Wait for the system to reach the quiet single-user mode)
.
# fsck /dev/dsk/your_root_file_system_device
```

Remember that the device file name for any mounted file system can be found using either the mount command or the bdf command. On a Series 300/400 this name is most likely /dev/dsk/0s0, on a 700 it is /dev/dsk/6s0 and on an 800 it might be /dev/dsk/c0d0s13.

10-9. WORKSESSION: Review Questions

1. What type of archive is the tar command used for?

Answer:

Simple backups of a few files or directories. It is too slow to use it as an effective system backup tool.

2. What is the main function of the fbackup command?

Answer:

To perform full and incremental system backups. It is very fast compared to other tools available on the system.

3. If you do not specify the f key to tar, where is the archive written?

Answer:

To the default device /dev/rmt/0m.

4. Can you append files or directories to an fbackup archive tape? What about a tar tape?

Answer:

You cannot append files to an fbackup archive.

Appending files to an archive is one of tar's best features.

10-10. LAB: Backing Up the System

1. Using the tar command, backup norton's home directory.

Answer:

```
# tar cvf device_name /users/norton
```

OR better:

```
# cd /users/norton
# tar cvf device_name .
```

device_name will be the name of your tape device. If your tape device is named /dev/rmt/0m you do not need the f *device_name* on the command line.

If there are not enough tape drives for every system in the class, you can specify a regular filename for the *device_name*.

By cd'ing to the directory to backup first then archiving ".", you can recover the files into any directory. Not just /users/norton.

2. List the contents of the tar tape.

Answer:

```
# tar tvf device_name
```

3. Append the file /etc/passwd to the tape.

Answer:

```
# tar rvf device_name /etc/passwd
```

4. Remove norton's .profile from his home directory then Recover *just* norton's .profile from the tape.

Answer:

```
# tar xvf device_name /users/norton/.profile
```

If you created the archive using "." instead of /users/norton you could go to any directory on the system and recover the .profile into that directory:

```
# cd /users/norton
# tar xvf device_name .profile
```

5. Interactively backup just the /users directory using SAM. Make sure you create a *disk* index file.

Answer:

On the Back Up Files Interactively... screen you must specify /users to be included in the backup. You could exclude someone's home directory from the backup by putting /users/*user_name* in the exclude column if you wanted to experiment more.

Answer y to the "Create index file?" question.

6. View the index file from the tape then view the index file from the disk. Are they the same?

Answer:

To get the index file from tape you must use Get List of Files from a Backup Tape ... and specify the correct device name. The index will be displayed on your screen and saved in the file you specified.

To see the index file created on disk during the backup select
View Information on Last: Full Backup

7. Make a new directory called /users/testdir then recover just norton's home directory into /users/testdir.

Answer:

On the "*Recover Files or Directories*" screen you need to "enter names of individual files or directories" and specify /users/norton in the include column.

The directory for recovery should be /users/testdir.

If you answered y to the full pathname question, you would recover /users/norton into /users/testdir. So you would actually see /users/testdir/users/norton/*norton's_files* in a listing.

Answering n to the full pathname question, would recover just *norton's_files* directly into /users/testdir.

11-11. WORKSESSION: Review Questions

1. What is the difference between a Series 700 recovery system and restore tape?

Answer:

The restore tape will completely rewrite the root disk whereas the recovery system will give you tools so you can attempt to fix it without destroying the current contents.

2. What is the purpose of the `/usr/lib/cron/cron.allow` file?

Answer:

This file contains the list of users that are allowed to use `crontab`.

3. Can cron be used to bring the system to a single user run-level, perform a system backup then reboot back into the default run-level?

Answer:

NO. Once the system is changed to a single-user run-level all processes including your job and the cron daemon are gone. The system would sit there in single-user mode until someone manually brought it back up.

4. Why would you want to remove a software subsystem from your kernel?

Answer:

If you loaded a software package like the X.25 communications software on your system, the kernel would be configured to support X.25. If you remove the X.25 software using `rmfn` you should also take it out of the kernel to reduce its size.

11-12. LAB: Lab Exercises

1. Schedule a job to run every 15 minutes that displays the date to your screen.

Answer:

The cronfile would look something like this:

```
0,15,30,45 * * * * date > /dev/your_tty_name
```

You can find *your_tty_name* using the `who` or `tty` commands.

You then submit your cronfile to cron using `crontab`:

$ crontab *cronfile_name*

2. Run SAM and enter the Kernel Configuration -> menu. Go to the View/Modify I/O Configuration ... screen and note the name of the driver used to support SCSI DAT tapes.

Answer:

The driver name is scsitape on Series 300/400/700 and tape2 on the Series 800.

3. Press Exit Task to return to the "*Kernel Configuration*" screen then select Modify Operating System Parameters ->. Try to find the maxusers and maxdsiz parameter. Use the Help facility to determine what they are used for. Look around at some of the other kernel tunable parameters and see if you can tell what they are used for.

Answer:

maxusers is on the "*Miscellaneous Parameters*" screen. It is used to scale many other system parameters. It has nothing to do with how many users can log into the system. maxusers is one of the most commonly changed kernel parameters.

maxdsiz is the maximum amount of data space a process can allocate. Applications that use large data sets (like 3D CAD or large database applications) often need a larger value for maxdsiz. This parameter is found on the "*Process Parameters*" screen.

4. Series 300/400/700 Only: If the hardware is available in your classroom, make a recovery system using the mkrs command. If time permits, boot the recovery system and try a simple task like removing the root password.

Answer:

See the "Series 300/400/700 Recovery System" section of this module for the details.