

# HP-UX System Administrator Manual

Volume 2

HP 9000 Series 300 Computers

HP Part Number 98594-90060



**Hewlett-Packard Company**

3404 East Harmony Road, Fort Collins, Colorado 80525

---

## Legal Notices

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

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

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

Copyright © Hewlett-Packard Company, 1988

This document contains information which is protected by copyright. All rights are reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

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

Copyright © AT&T, Inc. 1980, 1984

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

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

---

## Printing History

New editions of this manual will incorporate all material updated since the previous edition. Update packages may be issued between editions and contain replacement and additional pages to be merged into the manual by the user. Each updated page will be indicated by a revision date at the bottom of the page. Note that pages which are rearranged due to changes on a previous page are not considered revised.

The manual printing date and part number indicate its current edition. The printing date changes when a new edition is printed. (Minor corrections and updates which are incorporated at reprint do not cause the date to change.) The manual part number changes when extensive technical changes are incorporated.

December 1988 ... Edition 1. Documents the 6.5 release of HP-UX.

This edition of the *HP-UX System Administrator Manual* replaces part number 97597-90060 (First Edition; April 1988).



# Contents

---

## 1. Getting Started

Welcome . . . . .	1-1
What's in this Manual? . . . . .	1-2
Conventions Used in this Manual . . . . .	1-3
Using Other HP-UX Manuals . . . . .	1-5
The Administrator's Responsibilities . . . . .	1-6
Installing and Testing the Hardware . . . . .	1-6
Evaluating Users' Needs . . . . .	1-6
Installing the HP-UX Operating System . . . . .	1-6
Configuring the HP-UX kernel . . . . .	1-7
Allowing Users Access to the System . . . . .	1-7
Adding and Moving Peripheral Devices . . . . .	1-7
Monitoring File System Use and Growth . . . . .	1-7
Updating the HP-UX System . . . . .	1-7
Back Up, Recovery, and Restoring the System . . . . .	1-8
Detecting/Correcting File System Errors . . . . .	1-8
Assisting Other Users . . . . .	1-9
Providing a "Back-up" Administrator . . . . .	1-9
User Survey . . . . .	1-10

## 2. System Management Concepts

The HP-UX Hierarchy . . . . .	2-2
Processes . . . . .	2-6
Process Creation (Parent and Child Processes) . . . . .	2-6
Process Termination . . . . .	2-8
Process Groups . . . . .	2-8
Terminal Affiliation . . . . .	2-9
Open Files in a Process . . . . .	2-9
IDs . . . . .	2-10

The Superuser . . . . .	2-12
File System Implementation . . . . .	2-13
Disk Layout . . . . .	2-14
Data Storage . . . . .	2-20
Updating the HFS File System . . . . .	2-23
Corruption of the File System . . . . .	2-25
Detection and Correction of Corruption . . . . .	2-25
File Format and Compatibility . . . . .	2-31
File Protection . . . . .	2-32
File Sharing and Locking . . . . .	2-38
The File System Buffer Cache . . . . .	2-40
Magnetic Tape . . . . .	2-42
Magnetic Tape Definitions . . . . .	2-42
Preventive Maintenance . . . . .	2-47
Tape Streaming . . . . .	2-47
Optical Technology . . . . .	2-51
Rewritable Optical . . . . .	2-51
Why Use Rewritable Optical? . . . . .	2-52
Hewlett-Packard's Rewritable Optical Products . . . . .	2-53
HP Series 6300 Model 650/A - Optical Disk Drive . . . . .	2-53
Optical Disk Drive Guidelines . . . . .	2-54
HP Series 6300 Model 20GB/A - Optical Autochanger . . . . .	2-54
Optical Autochanger Guidelines . . . . .	2-55
Memory Management . . . . .	2-58
Overview . . . . .	2-58
Logical Address Space Management . . . . .	2-59
Physical Memory Utilization . . . . .	2-62
Swap Space Management . . . . .	2-63
Shared Code . . . . .	2-65
Demand Load . . . . .	2-67
Device I/O . . . . .	2-68
Device Classes . . . . .	2-68
Drivers . . . . .	2-68
Cluster Concepts . . . . .	2-69
Why Use a Cluster? . . . . .	2-69
What is an HP-UX Cluster? . . . . .	2-71
Model of Cluster System Administration . . . . .	2-74
Context . . . . .	2-77

Context Dependent Files (CDFs) . . . . .	2-78
CDF Examples . . . . .	2-85
Cnode-Specific Device Files . . . . .	2-90
Cluster Server Processes . . . . .	2-92
Process IDs . . . . .	2-94
Swapping in a Cluster . . . . .	2-95
Tips on Customizing your Cluster . . . . .	2-96
Changes to Commands and Files . . . . .	2-97
Subsystem Administration . . . . .	2-105
<b>3. System Startup and Shutdown</b>	
System Startup Functions . . . . .	3-1
Booting the System (Standalone or Root Server) . . . . .	3-2
Booting a Cluster . . . . .	3-5
Overview of Internal Functions of System Startup . . . . .	3-7
The Boot ROM . . . . .	3-8
HP-UX Takes Control . . . . .	3-15
HP-UX Starts the Init Process . . . . .	3-16
Init Brings the System to Run-Level 2 . . . . .	3-19
Init Spawns gettys to Cause a Login Prompt . . . . .	3-23
A User Logs In . . . . .	3-24
System Administration Mode . . . . .	3-26
Booting Problems . . . . .	3-27
Shutting Down the System . . . . .	3-28
Shutting Down for System Maintenance . . . . .	3-28
Halting or Rebooting the System . . . . .	3-29
Shutting Down a Cluster . . . . .	3-30
Power Fail or Disk Crash Recovery . . . . .	3-32
<b>4. Customizing the HP-UX System</b>	
Changing the HP-UX Environment Files . . . . .	4-3
/etc/inittab . . . . .	4-3
/etc/rc . . . . .	4-4
/etc/passwd . . . . .	4-4
/etc/group . . . . .	4-4
/etc/motd . . . . .	4-4
/usr/news . . . . .	4-5
/etc/profile or /etc/csh.login . . . . .	4-5

/etc/wtmp . . . . .	4-5
/etc/btmp . . . . .	4-6
/etc/utmp . . . . .	4-6
/etc/securetty . . . . .	4-6
\$HOME/.profile, \$HOME/.cshrc, \$HOME/.login, or \$HOME/.environ . . . . .	4-7
\$HOME/.exrc . . . . .	4-8
/usr/lib/terminfo . . . . .	4-8
/etc/checklist . . . . .	4-8
/etc/catman . . . . .	4-9
/etc/issue . . . . .	4-10
/etc/csh.login, /etc/rc, and /etc/profile . . . . .	4-10
/usr/lib/tztab . . . . .	4-11
/etc/ttytype . . . . .	4-11
Adding Peripheral Devices . . . . .	4-12
Overview of the Task . . . . .	4-12
Miscellaneous Devices . . . . .	4-18
Terminals and Modems . . . . .	4-19
Pseudo Terminals . . . . .	4-36
Hard Disks . . . . .	4-39
Rewritable Optical Disk Drives and Optical Autochangers	4-46
Flexible Disk Drive . . . . .	4-56
Cartridge Tape Drive . . . . .	4-61
Nine-Track Magnetic Tape . . . . .	4-63
Plotters and Digitizers . . . . .	4-66
HP-HIL Devices . . . . .	4-68
GPIO Devices . . . . .	4-68
Setting Up Data Communication on your System . . . . .	4-69
Setting the Terminal Characteristics . . . . .	4-71
Selecting a Value for the TERM Variable . . . . .	4-71
Setting TERM with the tset Command . . . . .	4-73
Adding to /etc/checklist . . . . .	4-74
Configuring your HP-UX Cluster . . . . .	4-78
Prerequisites and Assumptions . . . . .	4-78
HP-UX Cluster Information Sheet . . . . .	4-78
Creating the Clustered Environment . . . . .	4-83
Adding a Diskless Cluster Node . . . . .	4-92
Before Adding a Cnode . . . . .	4-92



Procedure . . . . .	4-92
After Adding a Cnode . . . . .	4-98
Removing and Renaming a Cluster Node . . . . .	4-102
Removing a Diskless Cluster Node . . . . .	4-102
Renaming a Diskless Cluster Node . . . . .	4-106
Creating a New File System . . . . .	4-107
Prerequisites . . . . .	4-107
Creating the File System . . . . .	4-108
After Creating the File System . . . . .	4-109
Enabling Long File Names . . . . .	4-113
Main Differences Between Long and Short File Names . . . . .	4-113
Long File Names or Short File Names? . . . . .	4-114
Enabling Long File Names on an Existing File System . . . . .	4-115
Long File Names and Your System . . . . .	4-118
Creating System Run-Levels . . . . .	4-125
Guidelines for Creating New System Run-Levels . . . . .	4-126
Example /etc/inittab File . . . . .	4-127
Creating and Using a Recovery System . . . . .	4-129
Creating a Recovery System . . . . .	4-130
Using the mkrs Script . . . . .	4-132
Booting and Using the Recovery System . . . . .	4-133
Shutting Down the Recovery System . . . . .	4-135
Recovering the System Manually . . . . .	4-135
Mounting and Unmounting File Systems . . . . .	4-140
To Mount a File System . . . . .	4-141
To Unmount a File System . . . . .	4-144
Mounting/Unmounting File Systems Using /etc/checklist . . . . .	4-145
Mounting the /usr File System . . . . .	4-146
Removing Optional Products and Filesets . . . . .	4-148
The reconfig Command . . . . .	4-149
Using reconfig . . . . .	4-150
Additional Notes on reconfig . . . . .	4-153
Setting Up Printers and the LP Spooler . . . . .	4-154
What Is in This Section . . . . .	4-155
LP Spooler Terminology and Overview . . . . .	4-156
Installing Printers and the LP Spooler . . . . .	4-157
Configuring the LP Spooler for Remote Operation . . . . .	4-168
General-Purpose LP Spooler Commands . . . . .	4-172

System Administrator LP Spooler Commands . . . . .	4-174
Remote Spooling Commands . . . . .	4-175
LP Spooling Directories . . . . .	4-176
LP Spooler Administrator Duties . . . . .	4-177
How Models Work . . . . .	4-181
Setting Up a Printer (the Manual Method) . . . . .	4-182
The LP Spooler in an HP-UX Cluster . . . . .	4-184
LP Spooler Errors . . . . .	4-184
Installing Optional Software and Updating Your HP-UX	
Kernel . . . . .	4-186
Update Process Overview . . . . .	4-187
Overview of Updating if You Will Use an HP-UX Cluster	4-188
General Discussion of the Update Process . . . . .	4-189
Preparing to Modify Your System . . . . .	4-191
Updating or Installing Optional Software from Cartridge	
Tape . . . . .	4-196
Updating or Installing Optional Software from Flexible	
Disk . . . . .	4-212

## 5. Periodic System Administrator Tasks

Adding/Removing Users . . . . .	5-3
Adding a New User (Using Reconfig) . . . . .	5-4
Adding a New User (the Manual Method) . . . . .	5-12
Creating the /etc/passwd Entry . . . . .	5-14
Setting the New User's Password . . . . .	5-16
Removing a User (Using Reconfig) . . . . .	5-17
Removing a User from the System (the Manual Method)	5-20
Suspending a User from the System . . . . .	5-20
Backing Up and Restoring the File System . . . . .	5-21
Background Information . . . . .	5-21
Backing Up your File System Using the Backup Scripts .	5-27
Backing Up the File System Using the Fbackup Command	5-32
Backing Up the File System Using the Optical Disk	
Autochanger . . . . .	5-36
Backing Up Selected Files onto Cartridge Tape . . . . .	5-41
Backing Up Selected Files onto Flexible Disk or Magnetic	
Tape . . . . .	5-42
Backing Up Selected File Onto Rewritable Optical Disks	5-43

Restoring the System . . . . .	5-43
Performing Backups Automatically . . . . .	5-47
Changing a Password . . . . .	5-50
Changing the System's Run-Level . . . . .	5-51
Entering Run-Level s (The System Administration Mode)	5-51
Changing Run-levels . . . . .	5-52
Run-Level Review . . . . .	5-53
Creating Groups/Changing Group Membership . . . . .	5-54
Communicating with System Users . . . . .	5-56
Controlling Disk Use . . . . .	5-58
Initializing Media . . . . .	5-61
What Needs to be Initialized? . . . . .	5-61
Prerequisites . . . . .	5-62
Using mediainit . . . . .	5-62
Initializing Media to LIF Format . . . . .	5-64
Setting the System Clock . . . . .	5-66
Setting the Time Zone . . . . .	5-66
Setting the Time and Date . . . . .	5-67
Possible Problems When Changing the System Clock . . . . .	5-68
Transferring Files . . . . .	5-70
cpio . . . . .	5-70
tcio . . . . .	5-71
tar . . . . .	5-72
dd . . . . .	5-73
LIF Utilities . . . . .	5-73
UUCP . . . . .	5-76
Local Area Network . . . . .	5-77
Virtual Terminal . . . . .	5-77
Moving Directories . . . . .	5-77

## 6. Kernel Customization

Introduction . . . . .	6-1
Using the config Command . . . . .	6-3
Files Required for the config Process . . . . .	6-3
Running config . . . . .	6-3
Configuring Device Drivers and I/O Cards . . . . .	6-8
Sample Configuration Description Files . . . . .	6-8
Determining the Required Device Drivers . . . . .	6-11

Configuring Swap Space . . . . .	6-13
Overview of Creating More Swap Space on an Additional Disk . . . . .	6-15
Step 1: Initialize the Media . . . . .	6-16
Step 2: Determine the Amount of Swap Space Needed . . . . .	6-16
Step 3: Create the File System on the New Disk . . . . .	6-22
Step 4: Add the Appropriate Entries into Dfile . . . . .	6-22
Step 5: Edit the /etc/checklist file . . . . .	6-23
Step 6: Execute the Config Command . . . . .	6-24
Step 7: Enable the New Swap Space if You Skipped Step 5 . . . . .	6-24
Possible Swap Setups . . . . .	6-24
Configuring Operating System Parameters . . . . .	6-29
HP-UX Operating System Parameters . . . . .	6-30
System V IPC Code . . . . .	6-31
Configuring Device Drivers and I/O Cards Using reconfig . . . . .	6-33
Creating the Minimum Operating System . . . . .	6-33
Creating a Fully Loaded Operating System . . . . .	6-38
Creating a Custom Operating System . . . . .	6-41
Creating a New Kernel in an HP-UX Cluster . . . . .	6-52
Local and Remote Swapping in an HP-UX Cluster . . . . .	6-56

## 7. System Accounting

What Is in This Chapter? . . . . .	7-2
Installation and Daily Usage . . . . .	7-3
How to Install System Accounting . . . . .	7-3
Summary of Daily Operation . . . . .	7-5
Overview of System Accounting . . . . .	7-7
Definitions . . . . .	7-7
Introduction to Commands . . . . .	7-9
System Data Flow . . . . .	7-13
Login and Directory Structure . . . . .	7-15
Disk Space Usage Accounting . . . . .	7-17
Reporting Disk Space Usage . . . . .	7-17
Creating Total Accounting Records . . . . .	7-21
Connect Session Accounting . . . . .	7-23
Writing Records to wttmp – acctwttmp . . . . .	7-23
Displaying Connect Session Records – fwttmp . . . . .	7-24

Fixing wtmp Errors – wtmpfix . . . . .	7-26
Creating Total Accounting Records . . . . .	7-26
Process Accounting . . . . .	7-30
Turning Process Accounting On . . . . .	7-30
Turning Process Accounting Off . . . . .	7-32
Checking the Size of pacct . . . . .	7-32
Displaying Process Accounting Records – acctcom . . . . .	7-35
Command Summary Report – acctems . . . . .	7-44
Creating Total Accounting Records . . . . .	7-48
Charging Fees to Users – chargefee . . . . .	7-50
Summarizing and Reporting Accounting Information . . . . .	7-51
Displaying Total Accounting Records – prtacct . . . . .	7-51
Merging Total Accounting Files – acctmerg . . . . .	7-54
Creating Daily Accounting Information – runacct . . . . .	7-58
Displaying runacct Reports – prdaily . . . . .	7-64
Creating Monthly Accounting Reports – monacct . . . . .	7-67
Updating the Holidays File . . . . .	7-68
Fixing Corrupted Files . . . . .	7-69
Fixing wtmp Errors . . . . .	7-69
Fixing tacct Errors . . . . .	7-70
Sample Accounting Shell Scripts . . . . .	7-71
grpdsug . . . . .	7-71
acct_bill . . . . .	7-73
System Accounting Files . . . . .	7-77
Files in the /usr/adm Directory . . . . .	7-77
Files in the /usr/adm/acct/nite Directory . . . . .	7-78
Files in the /usr/adm/acct/sum Directory . . . . .	7-79
Files in the /usr/adm/acct/fiscal Directory . . . . .	7-79

## 8. Trusted Systems

Components of System Security . . . . .	8-3
Physical Security . . . . .	8-3
Procedural Security . . . . .	8-4
System Security . . . . .	8-4
Security Policy . . . . .	8-5
Converting to a Trusted System . . . . .	8-7
Trusted Computing Base (TCB) . . . . .	8-9
Introduction to the TCB . . . . .	8-9

Parts of the TCB . . . . .	8-9
Access Control Lists (ACLs) . . . . .	8-10
Auditing . . . . .	8-11
Auditing Concepts . . . . .	8-12
Auditing Tasks . . . . .	8-20
Starting or Halting the Auditing System . . . . .	8-21
Selectively Auditing Users . . . . .	8-23
Auditing Selectively by Events . . . . .	8-24
Display the Auditing Information . . . . .	8-26
Audit Overflow Prevention and Handling . . . . .	8-28
Auditing in Subsystems and Applications . . . . .	8-30
The Password File . . . . .	8-32
Security Procedure Awareness . . . . .	8-33
Additional Changes to Your System to Improve Security . . . . .	8-33
System Administration Security Practices . . . . .	8-34
Performing General Administration Tasks in a Trusted System . . . . .	8-37

**A. Using the fsck Command**

Introduction . . . . .	A-1
fsck Errors That May Occur . . . . .	A-5
Initialization Phase Errors . . . . .	A-6
Phase 1 Errors: Check Blocks and Sizes . . . . .	A-9
Phase 1b: Rescan for More Dups . . . . .	A-13
Phase 2: Check Path Names . . . . .	A-13
Phase 3: Check Connectivity . . . . .	A-19
Phase 4: Check Reference Counts . . . . .	A-20
Phase 5: Check Cylinder Groups . . . . .	A-23
Phase 6: Salvage Cylinder Groups . . . . .	A-25
Cleanup . . . . .	A-25

**B. Troubleshooting**

System Run-Time Errors . . . . .	B-2
Kernel Errors . . . . .	B-2
File System Errors . . . . .	B-4
Error Messages . . . . .	B-5
Booting Diagnostic Errors . . . . .	B-5
Error Messages During Login . . . . .	B-6

Unable to Access a Peripheral Device . . . . .	B-8
Dealing with an Unresponsive Terminal . . . . .	B-9
Backup Recovery Errors . . . . .	B-10
Update Errors . . . . .	B-13
Errors During Update . . . . .	B-13
To check your tapes: . . . . .	B-16
HP-UX Cluster Troubleshooting . . . . .	B-17
Unsolved Problems . . . . .	B-17
Flowchart Format . . . . .	B-18
Diskless Boot Troubleshooting . . . . .	B-18
Possible Boot Daemon Error Messages . . . . .	B-43
Possible Panic, Retry, and Other System Messages on Diskless Cnodes . . . . .	B-46
System Messages Cont.) . . . . .	B-52
LAN Cable Break Detection Behavior . . . . .	B-52
Unable to Login or Rlogin to a Diskless Cnode . . . . .	B-55

### C. Partitions and Filesets

How Partitions and Filesets Work . . . . .	C-1
Dependent Filesets . . . . .	C-2
Mandatory vs. Non-mandatory Partitions . . . . .	C-2
Which Filesets Do I Want? . . . . .	C-3
Filesets and Their Functions . . . . .	C-4
Filesets, Dependencies, and Sizes . . . . .	C-9
Recovering Lost Files . . . . .	C-15
Recovering Files Required For Update . . . . .	C-16
How to List Files . . . . .	C-17
List of AXE System Files . . . . .	C-18
Partition SYS_CORE . . . . .	C-18
Partition TEXT . . . . .	C-24
Partition SYS_TOOLS . . . . .	C-26
Partition STARBASE . . . . .	C-28
Partition X11WINDOWS . . . . .	C-29
Partition WINDOW . . . . .	C-33
Partition PROG_LANGS . . . . .	C-34
Partition MISC_UTILS . . . . .	C-34
Partition NLS . . . . .	C-45
List of PE System Files . . . . .	C-48

Partition SYS_CORE . . . . .	C-48
Partition TEXT . . . . .	C-49
Partition SYS_TOOLS . . . . .	C-79
Partition STARBASE . . . . .	C-84
Partition X11WINDOWS . . . . .	C-89
Partition WINDOW . . . . .	C-95
Partition PROG_LANGS . . . . .	C-98
Partition MISC_UTILS . . . . .	C-102
Partition NLS . . . . .	C-104

#### D. Operating System Parameters

acctr resume . . . . .	D-2
acctsuspend . . . . .	D-4
argdevnblk . . . . .	D-6
dos_mem_byte . . . . .	D-7
dskless_cbufs . . . . .	D-8
dskless_fsbufs . . . . .	D-10
dskless_mbufs . . . . .	D-12
dskless_node . . . . .	D-14
dst . . . . .	D-15
filesizelimit . . . . .	D-16
fpa . . . . .	D-17
maxdsiz . . . . .	D-18
maxssiz . . . . .	D-19
maxswapchunks . . . . .	D-20
maxtsiz . . . . .	D-22
maxuprc . . . . .	D-23
maxusers . . . . .	D-25
mesg . . . . .	D-27
minswapchunks . . . . .	D-28
msgmap . . . . .	D-29
msgmax . . . . .	D-31
msgmnb . . . . .	D-32
msgmni . . . . .	D-33
msgseg . . . . .	D-34
msgssz . . . . .	D-36
msgtql . . . . .	D-38
nbuf . . . . .	D-39



ncallout . . . . .	D-41
ndilbuffers . . . . .	D-43
netmeminit . . . . .	D-44
netmemmax . . . . .	D-45
netmemthresh . . . . .	D-46
nfile . . . . .	D-48
nflocks . . . . .	D-49
ngcsp . . . . .	D-50
ninode . . . . .	D-52
nproc . . . . .	D-53
npty . . . . .	D-55
ntext . . . . .	D-56
num_cnodes . . . . .	D-57
num_lan_cards . . . . .	D-59
parity_option . . . . .	D-60
scroll_lines . . . . .	D-62
selftest_period . . . . .	D-63
sema . . . . .	D-64
semaem . . . . .	D-66
semmap . . . . .	D-68
semmni . . . . .	D-70
semmns . . . . .	D-71
semmnu . . . . .	D-72
semume . . . . .	D-74
semvmx . . . . .	D-76
server_node . . . . .	D-77
serving_array_size . . . . .	D-79
shmall . . . . .	D-81
shmbmk . . . . .	D-82
shmem . . . . .	D-83
shmmax . . . . .	D-84
shmmaxaddr . . . . .	D-85
shmmin . . . . .	D-86
shmmni . . . . .	D-87
shmseg . . . . .	D-88
timeslice . . . . .	D-89
timezone . . . . .	D-91
unlockable_mem . . . . .	D-92

using_array_size . . . . .	D-94
Swap Space Parameter Interaction . . . . .	D-96

**Index**

Command Options . . . . .	7-56
Example . . . . .	7-57
Creating Daily Accounting Information – runacct . . . . .	7-58
Files Processed by runacct . . . . .	7-58
The States of runacct . . . . .	7-59
Recovering from Failure . . . . .	7-61
Restarting runacct . . . . .	7-62
Daily Reports . . . . .	7-63
Displaying runacct Reports – prdaily . . . . .	7-64
Daily Line Usage Report . . . . .	7-65
Daily Resource Usage Report . . . . .	7-66
Daily and Monthly Command Summary . . . . .	7-66
Last Login . . . . .	7-66
Creating Monthly Accounting Reports – monacct . . . . .	7-67
Updating the Holidays File . . . . .	7-68
Fixing Corrupted Files . . . . .	7-69
Fixing wtmp Errors . . . . .	7-69
Fixing tacct Errors . . . . .	7-70
Sample Accounting Shell Scripts . . . . .	7-71
grpdusg . . . . .	7-71
The Shell Script . . . . .	7-72
acct_bill . . . . .	7-73
Examples . . . . .	7-73
Output Produced by acct_bill . . . . .	7-73
The Shell Script . . . . .	7-75
System Accounting Files . . . . .	7-77
Files in the /usr/adm Directory . . . . .	7-77
Files in the /usr/adm/acct/nite Directory . . . . .	7-78
Files in the /usr/adm/acct/sum Directory . . . . .	7-79
Files in the /usr/adm/acct/fiscal Directory . . . . .	7-79



# Contents

---

<b>7. System Accounting</b>	
What Is in This Chapter? . . . . .	7-2
Installation and Daily Usage . . . . .	7-3
How to Install System Accounting . . . . .	7-3
Update /etc/rc . . . . .	7-3
Create crontab Entries . . . . .	7-4
Set PATH for Accounting Commands . . . . .	7-5
Summary of Daily Operation . . . . .	7-5
Overview of System Accounting . . . . .	7-7
Definitions . . . . .	7-7
prime/non-prime connect time . . . . .	7-7
process accounting records . . . . .	7-8
total accounting records . . . . .	7-8
Introduction to Commands . . . . .	7-9
Installation . . . . .	7-9
Disk Space Usage Accounting . . . . .	7-10
Connect Session Accounting . . . . .	7-10
Process Accounting . . . . .	7-11
Charging Fees . . . . .	7-12
Summarizing and Reporting Accounting Information . . . . .	7-12
System Data Flow . . . . .	7-13
Login and Directory Structure . . . . .	7-15
Logging In . . . . .	7-15
Directory Structure . . . . .	7-15
Disk Space Usage Accounting . . . . .	7-17
Reporting Disk Space Usage . . . . .	7-17
acctdusg . . . . .	7-17
diskusg . . . . .	7-19
Creating Total Accounting Records . . . . .	7-21

acctdisk . . . . .	7-21
dodisk . . . . .	7-21
Connect Session Accounting . . . . .	7-23
Writing Records to wtmp – acctwtmp . . . . .	7-23
Displaying Connect Session Records – fwtmp . . . . .	7-24
Fixing wtmp Errors – wtmpfix . . . . .	7-26
Creating Total Accounting Records . . . . .	7-26
acctcon1 . . . . .	7-26
prctmp . . . . .	7-29
acctcon2 . . . . .	7-29
Process Accounting . . . . .	7-30
Turning Process Accounting On . . . . .	7-30
turnacct on . . . . .	7-31
accton . . . . .	7-31
Turning Process Accounting Off . . . . .	7-32
turnacct off . . . . .	7-32
accton . . . . .	7-32
Checking the Size of pacct . . . . .	7-32
ckpacct . . . . .	7-33
turnacct switch . . . . .	7-34
Displaying Process Accounting Records – acctcom . . . . .	7-35
Definitions of Information Produced by acctcom . . . . .	7-36
Report Format Options . . . . .	7-39
Record Selection Options . . . . .	7-40
Sample Reports . . . . .	7-43
Command Summary Report – acctcms . . . . .	7-44
Producing a Readable Report – the -a option . . . . .	7-44
Other Options . . . . .	7-46
Sample Report . . . . .	7-47
Creating Total Accounting Records . . . . .	7-48
acctprc1 . . . . .	7-49
acctprc2 . . . . .	7-49
Charging Fees to Users – chargefee . . . . .	7-50
Summarizing and Reporting Accounting Information . . . . .	7-51
Displaying Total Accounting Records – prtacct . . . . .	7-51
Report Format . . . . .	7-52
Example . . . . .	7-53
Merging Total Accounting Files – acctmerg . . . . .	7-54

## System Accounting

---

Multi-user HP-UX allows concurrent sharing of computer resources among multiple users: several users can be logged in, all sharing disk space, memory, and the CPU. On multi-user systems, HP-UX System Accounting provides the means to:

- Monitor disk space usage for individual users.
- Record connect session data (logins/logouts).
- Collect resource utilization data (such as memory usage and execution times) for individual processes.
- Charge fees to specific users.
- Generate summary files and reports that can be used to analyze system performance and bill users for resource consumption.

---

**Note**

Much of the material in this chapter assumes greater knowledge of HP-UX than is required of the “average” user. In particular, System Accounting borrows many concepts from two previous chapters—“System Management Concepts” and “System Startup and Shutdown”. If you are unfamiliar with the concepts and terminology in those chapters, then you should review them.

---

On an HP-UX cluster, you must enable/disable accounting on each cnode. Also, some concepts, such as disk space accounting, makes sense only on the root server.

---

## What Is in This Chapter?

HP-UX System Accounting allows you to accomplish accounting tasks through a number of versatile commands. This chapter illustrates the use of these commands and contains the following sections:

- “Installation and Daily Usage” shows the routine daily usage of System Accounting and shows you how to install it.
- “Overview of System Accounting” provides the background information necessary to understand how to use System Accounting.
- “Disk Space Usage Accounting” illustrates the use of the accounting commands that monitor disk space utilization on a per-user basis.
- “Connect Session Accounting” describes the commands that record and report connect session accounting information.
- “Process Accounting” shows how to generate per-process accounting data and reports.
- “Charging Fees to Users” is the section where you learn how to charge fees to users.
- “Summarizing and Reporting Accounting Information” shows how to generate the main daily and monthly accounting reports that are used to monitor system performance and bill users.
- “Updating the Holidays File” describes how to set up the file describing your holidays.
- “Fixing Corrupted Files” is useful when System Accounting files become inconsistent or messed up. This section discusses how to fix these files.
- “Sample Accounting Shell Scripts” provides listings of shell scripts that you might find useful on your system.
- “System Accounting Files” contains brief definitions of the files used by System Accounting.



---

## Installation and Daily Usage

The purpose of this section is to show you:

- What you must do to get System Accounting running on your system.
- How System Accounting automatically creates daily and monthly accounting data and reports.

After reading this section, you should be able to install System Accounting on your system. Once properly installed, System Accounting will automatically generate daily and monthly accounting data and reports.

### How to Install System Accounting

Not all users require accounting services on their systems. For this reason, HP-UX System Accounting is provided as an option: if you want to use System Accounting, you must install it yourself. The installation procedure is covered here.

There are three steps in the installation process:

1. Update `/etc/rc`
2. Create `crontab` entries
3. Set `PATH` for accounting commands

Each of these steps must be carried out to insure that System Accounting automatically creates daily and monthly accounting information. Detailed descriptions of each step follow.

#### Update `/etc/rc`

The system initialization shell script `rc` must be updated to automatically start System Accounting when the system is switched into multi-user mode. This requires adding the following entry in the `localrc` section of `/etc/rc`:

```
/bin/su - adm -c /usr/lib/acct/startup
```

## Create crontab Entries

To automate the daily and monthly creation of accounting data, you should create a `crontab` file that `cron` can use to automatically run certain accounting commands. This process entails the following steps:

1. Log in to System Accounting as the user `adm`.
2. Use an editor to create the `crontab` file containing the accounting commands that are to be run automatically by `cron`. (The actual entries to make in this file are shown after these steps.)
3. Execute the `crontab` command, specifying the file created in step 2 as input. This step insures that the `crontab` file created in step 2 will be scanned by `cron` every minute. After invoking this command, the step 2 file will be stored in the file:

```
/usr/spool/cron/crontabs/adm
```

4. At this point, you are finished creating `crontab` entries. If you ever want to change the entries, simply re-edit the file created in step 2 and use the `crontab` command again. Refer to the `crontab(1)` entry in the *HP-UX Reference* for more information.

The following entries, accompanied by a description of each, should be made in the `crontab` file created in above:

```
0 4 * * 1-6 /usr/lib/acct/runacct 2> /usr/adm/acct/nite/fd2log
```

`runacct`, the main accounting shell script, should be executed daily (during non-prime hours) to generate daily accounting reports. The above entry executes `runacct` at 4:00am every Monday through Saturday. Error messages will be redirected to the file `/usr/adm/acct/nite/fd2log`, if any errors occur while `runacct` executes.

```
0 2 * * 4 /usr/lib/acct/dodisk
```

`dodisk` creates total accounting records that summarize disk space usage for individual users. This entry runs `dodisk` at 2:00am every Thursday morning. *Do not add this entry if you are on a diskless node in an HP-UX cluster.*

```
5 * * * * /usr/lib/acct/ckpacct
```

To insure that the process accounting file, `pacct`, doesn't get too large, the command `ckpacct` should be executed hourly. This entry invokes `ckpacct` at five minutes into every hour.

```
15 5 1 * * /usr/lib/acct/monacct
```

The monthly merging of accounting data is facilitated through the `monacct` command. This entry allows `monacct` to generate a monthly total report and total accounting file. `monacct` will be executed at 5:15am on the first day of every month.

---

**Note** The dates and times shown in the `crontab` entries above are only suggestions; you can tailor `crontab` entries to suit your needs. However, if you use different entries than those shown here, be sure that `monacct` is run at such a time as to allow `runacct` sufficient time to finish.

---

## Set PATH for Accounting Commands

Finally, you should set the `PATH` shell variable in `/usr/adm/.profile` so that System Accounting knows where to look for commands. `Path` should be set as follows:

```
PATH=/usr/lib/acct:/bin:/usr/bin:/etc:/usr/adm
```

## Summary of Daily Operation

The daily operation of System Accounting is summarized by the following steps:

1. When HP-UX is switched into multi-user mode, the system initialization shell script `rc` executes the accounting command `startup`. The purpose of `startup` is to start System Accounting, and it performs the following functions:
  - a. Calls `acctwtmp` to add a boot record to `wtmp`. This record is marked by storing "acctg on" in the device name field of the `wtmp` record.
  - b. Turns process accounting on via `turnacct on`. `turnacct on` executes `accton` with the filename argument `/usr/adm/pacct`.
  - c. Removes work files left in the `sum` directory by `runacct`.

2. A report of the previous day's accounting information can be created by running `prdaily`. Obviously, this step is omitted the first day that System Accounting is installed, because the previous day's accounting information doesn't exist yet. However, after `runacct` has been executed, `prdaily` will generate valid reports.
3. The `ckpacct` command is executed every hour via `cron` to insure that the process accounting file `pacct` doesn't become too large. If `pacct` grows past a set maximum number of blocks, `turnacct switch` is invoked, which creates a new `pacct` file. (Other conditions may also limit the size of the process accounting file or turn process accounting off; for more details, refer to the discussion of `ckpacct` in the "Process Accounting" section of this chapter.) The advantage of having several smaller `pacct` files is that `runacct` can be restarted faster if a failure occurs while processing these records.
4. The `chargefee` program can be used to charge fees to users. It adds records to the file `fee`. These records are processed during the next execution of `runacct` and merged in with total accounting records.
5. `runacct` is executed via `cron` each night. It processes the active fee file and the process, connect session, and disk total accounting files. It produces command and resource-usage summaries by login name.
6. When the system is turned off using `shutdown`, the `shutacct` command is executed. The purpose of `shutacct` is to stop System Accounting, and it performs the following functions:
  - a. Writes a termination record to `wtmp` via the command `acctwtmp`. This record is marked by having "acctg off" in the device name field.
  - b. Turns process accounting off by calling `turnacct off`.

---

## Overview of System Accounting

In this section, the intrinsics of System Accounting are examined. Key terms are defined, commands are introduced, system data flow is described, and finally, you are shown the login and directory structure of System Accounting.

### Definitions

The following terms are specific to System Accounting. Other terms are listed in the Glossary at the end of this manual.

#### **prime/non-prime connect time**

Prime time is the time during the day when the computer system is most heavily used—for example, from 9:00am to 5:00pm. Non-prime time is the remaining time during the day when the system is less heavily used—from 5:00pm to 9:00am in this example.

When reporting computer time usage, System Accounting distinguishes between prime and non-prime time usage. You can specify prime and non-prime time on your system by editing the file `/usr/lib/acct/holidays`. (For details on the `holidays` file, refer to the section “Updating the Holidays File” in this chapter.

---

<b>Note</b>	Prime time is in effect only on weekdays (Monday through Friday); non-prime time is in effect during the weekends (Saturdays and Sundays) and on any holidays specified in the <code>holidays</code> file.
-------------	--

---

## **process accounting records**

Once System Accounting is installed and turned on, the following occurs: whenever a process terminates, the kernel writes a process accounting record for the terminating process into the current process accounting file, `/usr/adm/pacct` by default. (You can specify that a file other than `pacct` be used as the process accounting file, if desired.)

A process accounting record contains resource-usage data for a single process; it summarizes *how much* of the various resources the process used during its lifetime. Examples of information contained in process accounting records are:

- the user ID of the process's owner
- the name of the command that spawned the process
- the amount of time it took the process to execute

For greater detail on the contents and format of process accounting records, refer to *acct(4)* in the *HP-UX Reference Manual*.

## **total accounting records**

These records, created by various accounting commands, contain summary accounting information for individual users. These records provide the basic information for many reports generated by System Accounting. Some examples of information contained in these records are:

- the ID and user name of the user for whom the total accounting record was created
- the total number of processes that the user has spawned during the accounting period for which the total accounting record was created
- fees for special services rendered to this user

The exact contents and format of total accounting records can be found in *acct(4)*. In addition, commands covered in later sections of this chapter show how these records are created and used by System Accounting.

## Introduction to Commands

System Accounting provides many versatile commands to accomplish numerous, varied tasks. There are commands that create data, commands that display data, commands that remove data, commands that merge data, and commands that summarize and report data. In addition, the output of one command may become the input to other commands.

System Accounting commands can be logically categorized into six basic command groups:

- installation
- disk usage accounting
- connect session accounting
- process accounting
- charging fees
- summarizing and reporting accounting information

Descriptions of these command groups, along with a brief synopsis of each command, follow.

### Installation

These commands insure that System Accounting is properly installed. They are used to turn accounting on when HP-UX is powered up and turn accounting off when the system is shut down. They may also do some file cleanups. Two such commands exist:

- **startup**—starts accounting when HP-UX is switched to multi-user mode. **startup** is invoked from `/etc/rc`.
- **shutacct**—turns off accounting when HP-UX is turned off via the `/etc/shutdown` shell.

## Disk Space Usage Accounting

In general, these commands produce disk usage accounting information: they show disk space usage (in blocks) for individual users. They also produce total accounting records. There are four commands:

- `acctdusg` and `diskusg`—both commands show how many blocks of disk space users are consuming. They differ in command options, and the manner in which they produce the information—`acctdusg` takes its input from a list of path names created by `find`, and `diskusg` looks at the inodes of the file system to create its output.
- `acctdisk`—produces total accounting records. Its input is supplied (either directly or indirectly) from `acctdusg` or `diskusg`.
- `dodisk`—produces total accounting records by using the `diskusg` and `acctdisk` commands. `dodisk` is normally invoked by `cron`.

## Connect Session Accounting

Independently of System Accounting, the programs `login` and `init` record connect sessions by writing records into `/etc/wtmp`. System Accounting commands can display or fix this file, and can produce total accounting records for this file. There are six commands:

- `acctwtmp`—writes records to `wtmp`.
- `fwtmp`—displays the information contained in `wtmp`.
- `wtmpfix`—normalizes connect session records that span date changes (refer to *date(1)*). Also validates login names in connect session records.
- `acctcon1`—summarizes `wtmp` in ASCII readable format, producing one line per connect session.
- `acctcon2`—takes input of the format produced by `acctcon1` and produces total accounting records as output.
- `prctmp`—displays the session record file, normally called:

`/usr/adm/acct/nite/ctmp`



## Process Accounting

When process accounting is turned on, the kernel writes a process accounting record to `pacct` whenever a process terminates. A number of accounting commands exist that summarize and report this accounting information. In addition, certain commands turn process accounting on or off and insure that `pacct` doesn't become too large. The process accounting commands are:

- `accton`—turns process accounting on or off, depending on whether or not a filename argument is supplied with the command. If no filename is given, then process accounting is turned off; the kernel stops writing process accounting records to `pacct`. If a filename is specified, then the kernel starts writing process accounting records to the specified filename.  
  
`accton` uses the system call `acct` to turn process accounting on or off. Only the superuser can execute `accton`.
- `ckpacct`—checks the size of the process accounting file `pacct`. If `pacct` becomes too large, then a new `pacct` file is created via `turnacct switch`. *If disk space becomes critically short, then process accounting is turned off until sufficient space is available.* This command is normally invoked by `cron`.
- `turnacct on || off || switch`—performs one of three functions, depending on which argument (`on`, `off`, or `switch`) is specified. `turnacct on` turns process accounting on by calling `accton` with the default filename argument `/usr/adm/pacct`; `turnacct off` turns process accounting off by calling `accton` with no filename argument; `turnacct switch` renames the current `pacct` file (so that it is no longer the current process accounting file) and creates a new, empty `pacct` file.
- `acctcom`—displays process accounting records contained in `pacct` (or any specified file).
- `acctcms`—takes `pacct` as input, and produces summary accounting information by command, as opposed to by process.
- `acctprc1`—produces readable process accounting information, mainly for input into `acctprc2`.
- `acctprc2`—takes input of the form produced by `acctprc1` and produces total accounting records.

## Charging Fees

Occasionally, you may want to charge a user for something. For example, you might charge fees to users for fixing any damaged files that they have. The `chargefee` command allows you to charge fees to specific users.

## Summarizing and Reporting Accounting Information

This group of commands summarizes and reports the data created through the command groups described above. These are the commands that are probably used most frequently; they represent the highest level of accounting commands. Five such commands exist:

- `prtacct`—takes as input total accounting records and displays the records in ASCII readable format.
- `acctmerg`—combines the contents of separate total accounting files into a single total accounting file. This command allows the merging of disk, process, and connect session total accounting records.
- `runacct`—is the main accounting shell script. Normally invoked daily by `cron`, this command processes disk, connect session, process, and fee accounting information and produces summary files and reports. It accomplishes its task by proceeding through various states. In each successive state it invokes accounting commands to perform a specific task. For example, in one state, total accounting records for connect sessions are created; in another, disk, connect session, process, and fee total accounting records are merged to create one total accounting file.
- `prdaily`—invoked by `runacct` to format a report of the previous day's accounting data; the report is in the file `/user/adm/acct/sum/rptmmdd` where *mmdd* is the month and day of the report. `runacct` may also be used to display a report of the current day's accounting information.
- `monacct`—invoked once a month (or accounting period), this command summarizes daily accounting files and produces a summary files for the accounting period.

## System Data Flow

At this point, you have the rudimentary knowledge necessary to understand how System Accounting works; you know some important definitions and should basically know what the various commands do. The purpose of this section is to help you visualize how the different commands work together to create accounting data.

Figure 7-1 illustrates how accounting data is created. The diagram is broken into five separate sub-diagrams, each one representing the data flow for a given command group. The following notational conventions are used:

<i>Symbol</i>	<i>Description</i>
source $\Rightarrow$ dest	Wide arrows represent the transfer of data from a source to a destination. The source is at the start of the arrow; the destination, at the point. For example, the inodes of the file system are the source of information used by <code>diskusg</code> , which in turn is the source of disk usage reports that are inputs to <code>acctdisk</code> .
cause $\rightarrow$ object	Thin arrows represent cause-effect relationships. The cause lies at the start of the arrow; the object affected lies at the point. For example, <code>turnacct</code> on invokes <code>accton</code> which then signals the kernel to begin writing process accounting records to <code>pacct</code> .
files	Boxes with rounded corners represent files or groups of files. In a more general sense, they represent the inputs to and outputs from the various commands.

---


**Note** The installation commands do not appear in the diagram, because they aren't directly involved in the data creation process; they merely insure that it happens.

---

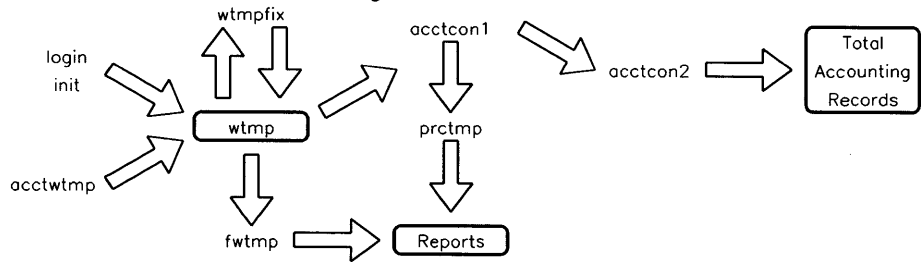
**Note** The commands `runacct` and `prdaily` are shown as having no inputs. This isn't exactly true: they do have inputs, but they get their inputs by executing other accounting commands. In essence, their inputs are the same basic inputs of the other command groups.

---

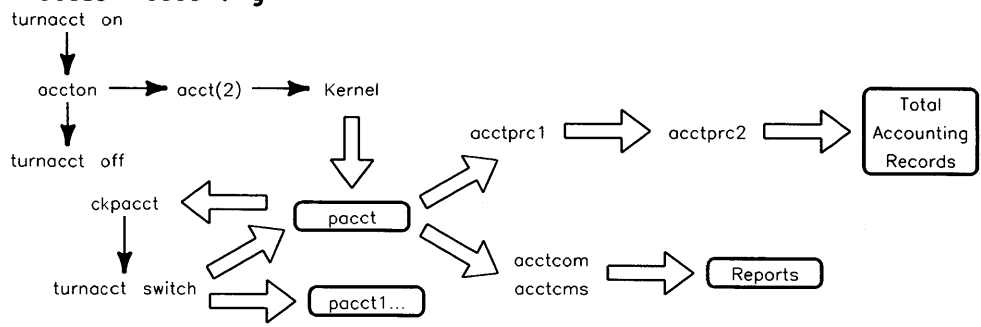
### Disk Usage Accounting




### Connect Session Accounting



### Process Accounting



### Charging Fees



### Summarizing and Reporting




Figure 7-1. System Accounting Data Flow Diagram

## Login and Directory Structure

You now know the basics, but you still can't begin learning the day-to-day usage of accounting commands until you know where to log in. In addition, you should know the accounting directory structure—where the various commands, directories, and files are located. These two topics are discussed here.

### Logging In

The login name for System Accounting is `adm`; the user ID for `adm` is 4. The `adm` login is a member of the *group* `adm`, and the group `adm` has a group ID of 4, also.

The home directory for the `adm` login is `/usr/adm`. You log in to System Accounting the same way you do for any account—simply supply the login name to the HP-UX login prompt:

```
login: adm
```

---

**Note** The integrity of accounting data files must be maintained if System Accounting is to generate accurate reports. For this reason, it is highly recommended that a password be used with the `adm` login.

---

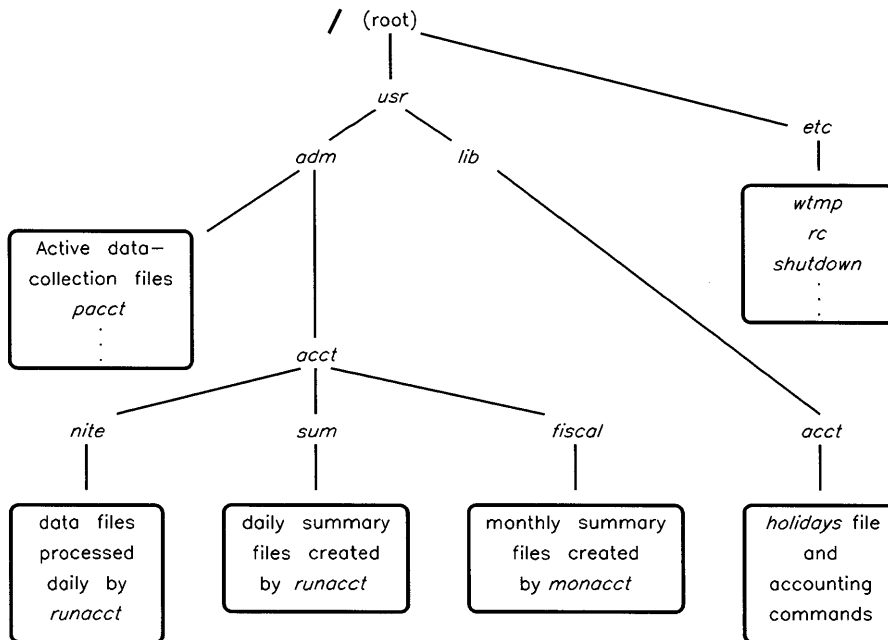
### Directory Structure

System Accounting uses a multi-level directory structure to organize its many accounting files. Each directory in this structure stores related groups of files, commands, or other directories. (Refer to the section “System Accounting Files” in this chapter for definitions of the accounting data files.)

Figure 7–2 illustrates this structure, and descriptions of each directory follow:

- `/usr/adm`—contains all active data-collection files, such as `pacct` and `fee`.
- `/usr/adm/acct`—contains the `nite`, `sum`, and `fiscal` directories described below.
- `/usr/adm/acct/nite`—stores data files that are processed daily by `runacct`.
- `/usr/adm/acct/sum`—cumulative summary files updated by `runacct` are kept here.

- /usr/adm/acct/fiscal—periodic (monthly) summary files created by monacct are stored here.
- /usr/lib/acct—System Accounting commands reside here.
- /etc—contains wtmp, and shell scripts rc and shutdown.



**Figure 7-2. System Accounting Directory Structure**

---

## Disk Space Usage Accounting

System Accounting provides the means to monitor disk space utilization for individual users. In this section, disk space usage accounting commands are explained. Before reading this discussion, you may want to review the “File System Implementation” section of Chapter 2.

Disk usage commands provide two main functions: they report disk usage (in blocks) for individual users and create disk total accounting records (supplied as inputs to commands such as `prtacct` or `runacct`).

### Reporting Disk Space Usage

Two commands—`acctdusg` and `diskusg`—report disk usage for individual users; both commands show the number of disk blocks allocated to specific users. However, each command has slightly different options. In addition, each differs in the manner in which it produces accounting information.

#### **acctdusg**

`acctdusg` takes from standard input a list of path names, usually created by the `find` command. For each file in the list, `acctdusg` identifies the owner of the file, computes the number of blocks allocated to the file, and adds this amount to a running total for the file’s owner. When finished looking through the list, `acctdusg` displays the information accumulated for each user: user ID, user name, and number of blocks used. Do not run this command on a diskless node in an HP-UX cluster; since the diskless nodes do not have a file system attached, the command does not provide any information.

This command is useful for reporting disk usage information for specific users or files. For example, suppose you want to know how many blocks of disk space you are using: your user ID is 351, user name is `bill`, and your home directory is `/users/pseudo/bill`. The following illustrates how you would use the `find` and `acctdusg` commands to show this information.

```
$ find /users/pseudo/bill -hidden -print > bills.files
$ acctdusg < bills.files
00351  bill                30
$ rm bills.files
```

In the above example, `bill` is using 30 blocks of disk space. The series of commands shown could easily have been combined into one line, such as:

```
$ find $HOME -hidden -print | acctdusg
00351  bill                30
```

The next example shows how to use `acctdusg` to generate disk usage information for all files in the system:

```
$ find / -hidden -print | acctdusg
00350  fred                  11
00351  bill                   30
00352  mike                   17
00353  sarah                   13
00354  molly                  18
00000  root                    3
00004  adm                    36
00001  bin                   2434
```

Two options are included with `acctdusg`:

- `-u no_owners` If `-u` is given, then path names of the files for which no owner is found are written into the file `no_owners`. This option could potentially find users who are trying to avoid disk charges.
- `-p p_file` The password file `/etc/passwd` is the default file used by `acctdusg` to determine ownership of files. If the `-p` option is used, then `acctdusg` will use `p_file` instead. This option is not needed if your password file is `/etc/passwd`.

The shell script `grpdsug`, provided in the section “Sample Accounting Shell Scripts” later in this chapter, displays disk accounting information for users in a given group. It illustrates the use of the `-u` option with `acctdusg`.



## diskusg

This command reports disk usage information in the same format as `acctdusg`—user ID, user name, and total disk blocks used. However, `diskusg` generates disk accounting information by looking through the *inodes* of a specified special file. (Refer to *inode(4)* and the “File System Implementation” section of Chapter 2 for more information on inodes and special files.) Therefore, `diskusg` is faster and more accurate than `acctdusg`. Do not run this command on a diskless cnode in an HP-UX cluster; since the diskless cnodes do not have a file system attached, the command does not provide any information.

The syntax of the `diskusg` command is:

```
diskusg [options] [files]
```

It generates a disk usage report from data in *files*, if specified; otherwise standard input is used. `diskusg` is normally invoked with the *files* argument. When specified, *files* are the special file names of the devices containing the inode information used by `diskusg` to generate its report. *files* is normally a special file from the `/dev` directory.

The following options may be used with `diskusg`:

- s This tells `diskusg` that: (1) input is in `diskusg` output format, and (2) that all lines for a single user should be combined into a single line. This option is used to merge data from separate files, each containing the output from using `diskusg` on different devices.
- v This option is useful for finding users who are trying to avoid disk space accounting charges. When this option is specified, `diskusg` writes records to `stderr` (standard error output) showing the special file name, inode number, and user ID of files that apparently have no owner.
- i *fnmlist* Causes `diskusg` to ignore the data on those file systems whose file system name is in *fnmlist*. *fnmlist* is a list of file systems separated by commas or enclosed within quotes.
- p *p\_file* This is the same as the `-p` option of `acctdusg`.
- u *u\_file* This option produces *exactly* the same output as the `-v` option. The difference between the two options is that `-v` writes its output to `stderr`; this option writes its output to the file *u\_file*.

The output of `diskusg` is normally used by `acctdisk` to create disk total accounting records. In addition, `diskusg` is normally called by `dodisk`.

The following example creates disk usage information for all users whose files reside on the disk whose device file is `/dev/rdisk/0s0`. *Note that the file system used in this example is the same as was used in the previous `acctdusg` example.*

```
$ diskusg /dev/rdisk/0s0
0      root      10616
1      bin       778
4      adm       96
350    fred      14
351    bill      32
352    mike      20
353    sarah     16
354    molly     22
355    julie     2
501    guest     2
```

The differences between `diskusg` and `acctdusg` are best illustrated by comparing their outputs. Note that:

1. `acctdusg` places leading zeros on user IDs; `diskusg` doesn't.
2. `acctdusg` counts files *only under each users \$HOME directory*. Files that users own in directories other than their home directory (for example, files in the `/tmp` directory) are counted as files with no owner.
3. Two extra users—`julie` and `guest`—show up in the output of `diskusg` when compared with the output from `acctdusg`. This occurred because the directories of these two users were empty; therefore, no disk usage totals were generated by `acctdusg`. However, `diskusg` looked at inodes and saw that `julie` and `guest` were actually using two blocks for the directories themselves.
4. If two or more users have links to a particular file, then `acctdusg` will prorate disk space usage for the file between each user. For example, if three users had a link to a 300-block file called `skurbnich.dat`, each user would be charged for 100 blocks of this file.

## Creating Total Accounting Records

Two commands are used to create total accounting records: `acctdisk`, and `dodisk`.

### **acctdisk**

`acctdisk` uses standard input records of the format produced by `acctdusg` and `diskusg`. From these records, `acctdisk` produces disk total accounting records that may be inputs to `prtacct` or `runacct`.

The following would write disk total accounting records to the file `disktacct` for all users in the group `pseudo`:

```
find / -group pseudo -print | acctdusg | acctdisk > disktacct
```

The next example would generate disk total accounting records for all users who have files on the disk `/dev/rdisk/1s0`. The total accounting records are written to the file `disktacct`.

```
diskusg /dev/rdisk/0s0 | acctdisk > disktacct
```

`acctdisk` has no options and is normally invoked by `dodisk`.

### **dodisk**

`dodisk` is normally invoked by `cron` to create disk total accounting records for daily usage by System Accounting. The syntax for `dodisk` is:

```
dodisk [-o] [files ...]
```

In the default case, `dodisk` creates disk total accounting records on the special files whose names are stored in `/etc/checklist`; the special file names are supplied as input to `diskusg`, which pipes its output to `acctdisk`, which in turn creates total accounting records.

If the `-o` option is used, `dodisk` creates total accounting records more slowly by using `acctdusg` instead of `diskusg`.

If *files* are used, disk accounting will be done on these file systems only. When the `-o` option is used, *files* should be mount points of mounted file systems; if omitted, *files* should be the special file names of mountable file systems.

---

**Note** Refer to the “Installation and Daily Usage” section of this chapter for more information on how `dodisk` should be invoked by `cron`.

---

It is possible for malicious users to defeat disk space accounting by giving their files away to other users with `chown(2)` or `chown(1)` (by default, all users can execute them). To avoid this, take away the ability to use these commands from some or all users with the `setprivgrp(1M)` command. To let only the superuser use the change-ownership abilities, add the following line to `/etc/rc`:

```
setprivgrp -n CHOWN
```

To let one or more groups of users use the change-ownership abilities, add a line for each group to `/etc/rc`, similar to the following:

```
setprivgrp group_name CHOWN
```

---

**Note** Taking away the change-ownership ability may cause problems when running some commands or applications.

---

---

**Note** In an HP-UX cluster, all cnodes normally share the same section of `/etc/rc` that sets privileged groups, and so all nodes have the same behavior with respect to privileged groups. However, `/etc/rc` can be coded such that different cnodes have different privilege group permissions in effect.

---

---

## Connect Session Accounting

Whenever a user logs in or out of HP-UX, the program `login` records the connect session in the file `/etc/wtmp`. Records in `wtmp` contain the following information:

- the terminal name on which the connect session occurred
- the login name of the user
- the current time/date at login or logout
- other status information (refer to *utmp(4)* for details)

System Accounting provides commands that allow you to write records to `wtmp`, to display and manipulate `wtmp`, and to create total accounting records from `wtmp`. These commands are covered in this section.

### Writing Records to `wtmp` – `acctwtmp`

The command `acctwtmp` allows you to write records to `wtmp` for whatever reason you might have. `acctwtmp` is normally invoked by `startup` and `shutacct` to record when System Accounting was turned on and off, respectively. The format of the command is:

```
acctwtmp 'reason'
```

where *reason* is a string describing the reason for writing the record to `wtmp`. Note that `acctwtmp` does not directly write records to `wtmp`: it writes a record containing the terminal name, current time, and reason string to standard output. To actually write the record to `wtmp` you must append the output from `acctwtmp` to the `wtmp` file as follows:

```
acctwtmp 'reason' >>/etc/wtmp
```

The *reason* string may be any combination of letters, numbers, spaces, and the dollar sign (\$), but may not exceed 11 characters in length. (*reason* must be enclosed in double quotes as shown.)

## Displaying Connect Session Records – `fwtmp`

To display the contents of `wtmp`, you can use the command `fwtmp`. When no options are used, `fwtmp` uses standard input records of the format contained in `wtmp`; it writes to standard output the ASCII readable equivalent of the input records. Two alternatives exist for the output from this command:

- The output of this command can be edited, via an HP-UX editor such as `vi`, and then rewritten to `wtmp` using special `fwtmp` options described below.
- The output can be supplied as input to commands which convert the information to total accounting records.

The syntax of `fwtmp` is:

```
fwtmp [-ic]
```

If no option is specified for the `fwtmp` command, then input is in binary format and is to be converted to ASCII readable format. The various combinations of the options `i` and `C` provide other combinations of input and output formats. The possible options are described below:

<b>Option</b>	<b>Description</b>
<code>-ic</code>	Input is in ASCII readable form and is to be converted to binary form. This is essentially the opposite of using <code>fwtmp</code> without any options.
<code>-i</code>	Both input and output are in ASCII readable format. This is the same as performing an ASCII to ASCII copy.
<code>-c</code>	Both input and output are in binary format—a binary to binary copy.

The following example shows the output produced by `fwtmp`. It is followed by a description of each column in the report:

```
$ fwtmp < /etc/wtmp
      system boot      0 2 0000 0000 479472540 Mar 12 03:49:00 1985
root   co console      0 7 0000 0000 479475173 Mar 12 04:32:53 1985
      acctg on         0 9 0000 0000 479493135 Mar 12 09:32:15 1985
mike   a1 ttya1        352 7 0000 0000 479493590 Mar 12 09:40:00 1985
mike   a1 ttya1        352 8 0011 0000 479496000 Mar 12 10:20:00 1985
sarah  07 tty07         353 7 0000 0000 479518335 Mar 12 16:32:15 1985
bill   10 tty10         351 7 0000 0000 479521475 Mar 12 17:24:35 1985
sarah  07 tty07         353 8 0011 0000 479522478 Mar 12 17:41:18 1985
bill   10 tty10         351 8 0011 0000 479526487 Mar 12 18:48:07 1985
      co console      0 8 0011 0000 479526488 Mar 12 18:48:08 1985
      acctg off       0 9 0000 0000 479526493 Mar 12 18:48:13 1985
      system boot      0 2 0000 0000 479389800 Mar 12 05:00:00 1985
```

## Column

## Description

- 1 The login name of the user who logged in or out.
- 2 `/etc/inittab` ID (this is usually the number of the line on which the connect session took place).
- 3 The name of the device on which the connect session occurred.
- 4 Process ID of the user who logged in or out.
- 5 Entry type. This field contains information on the type of record—for example, it shows whether the record is a login record (entry type=7), logout record (entry type=8), or if the record was written by `acctwtmp` (entry type=9). Refer to *utmp(4)* for more details on this field.
- 6-7 Exit status for connect session. Refer to *login(1)* and *utmp(4)* for details.
- 8 Time that entry was made (in elapsed seconds since January 1, 1970).
- 9-12 The equivalent of column 8 in date/time format showing month, day, time of day (in 24-hour format), and year.

## Fixing wtmp Errors – wtmpfix

When a user logs into HP-UX, the `login` program stores the value seven (7) in the entry type field of the connect session record. When the same user logs out, an entry type of eight (8) is recorded. You can see this by examining the sample output created by `fwtmp` in the previous section. Note that in the example, login records precede their corresponding logout records in chronological order.

Occasionally, this time-stamped ordering becomes inconsistent: logout records might precede login records. (This occurs when the date and time are reset while users are still logged in.) When this happens, the commands that create connect session total accounting records will not work properly.

Fortunately, the command `wtmpfix` fixes corrupted `wtmp` files. `wtmpfix` takes `wtmp` binary records as input and corrects the time/date stamps to be consistent; its standard output is also binary `wtmp` records. The syntax for `wtmpfix` is:

```
wtmpfix [files]
```

If *files* is given, then input is taken from *files*. A dash (-) can be used in place of *files* to indicate standard input. *Note that if you specify wtmp as both input to and output from this command, wtmp will be destroyed.* Therefore, take care not to destroy `wtmp`. The following shows how to properly fix `wtmp` using `wtmpfix`:

```
$ wtmpfix /etc/wtmp > wtmp.temp
$ fwtmp -c < wtmp.temp > /etc/wtmp
$ rm wtmp.temp
```

## Creating Total Accounting Records

This final set of connect session accounting commands is used to create connect session total accounting records. Before reading any further, you may want to review Figure 7-1 (in the “System Data Flow” section of this chapter).

### **acctcon1**

`acctcon1` converts a sequence of login/logoff records (of the format contained in `wtmp`) read from its standard input to a sequence of records, one per login session. Its input is normally redirected from `wtmp`; its output is columnar ASCII and can be supplied as input to `prctmp` or `acctcon2`.



The use of `acctcon1` is illustrated below by first displaying the contents of `wtmp` with `fwtmp`, and then using `acctcon1` to create a connect session summary file. The columnar data produced by `acctcon1` is described after the report:

```

$ fwtmp < /etc/wtmp
      system boot    0 2 0000 0000 479472540 Mar 12 03:49:00 1985
root   co console    0 7 0000 0000 479475173 Mar 12 04:32:53 1985
      acctg on       0 9 0000 0000 479493135 Mar 12 09:32:15 1985
mike   a1 ttya1      352 7 0000 0000 479493590 Mar 12 09:40:00 1985
mike   a1 ttya1      352 8 0011 0000 479496000 Mar 12 10:20:00 1985
sarah  07 tty07      353 7 0000 0000 479518335 Mar 12 16:32:15 1985
bill   10 tty10      351 7 0000 0000 479521475 Mar 12 17:24:35 1985
sarah  07 tty07      353 8 0011 0000 479522478 Mar 12 17:41:18 1985
bill   10 tty10      351 8 0011 0000 479526487 Mar 12 18:48:07 1985
      co console    0 8 0011 0000 479526488 Mar 12 18:48:08 1985
      acctg off      0 9 0000 0000 479526493 Mar 12 18:48:13 1985
$ acctcon1 < /etc/wtmp
20095488 353 sarah 1665 2478 479518335 Tue Mar 12 16:32:15 1985
521012224 352 mike          479493590 Tue Mar 12 09:40:00 1985
520095488 351 bill 0 5012 479521475 Tue Mar 12 17:24:35 1985
521011712 0 root 41047 6488 479475173 Tue Mar 12 04:32:53 1985

```

**Column**

**Description**

- 1 Shows the device address (in decimal equivalent of major/minor device address) at which the connect session occurred.
- 2 Gives the user ID for the connect session record.
- 3 Displays the login name for the user.
- 4 Shows the number of prime connect time seconds that were used during the connect session.
- 5 Shows non-prime connect seconds.
- 6 Displays the connect session starting time (in seconds elapsed since January 1, 1970).
- 7-11 Shows the conversion of column six to date/time format showing month, day time of day (in 24-hour format), and year.

In addition to its normal usage, `acctcon1` has four options:

Option	Description
<code>-p</code>	This option tells <code>acctcon1</code> not to produce one record per connect session. Instead, <code>acctcon1</code> simply echoes its input—one line per <code>wtmp</code> record—showing line name, login name, and time (in both seconds and day/time format). Using this option is similar to using <code>fwtmp</code> , except that this option doesn't show status information, whereas <code>fwtmp</code> does.
<code>-t</code>	<code>acctcon1</code> maintains a list of lines on which users are logged in. When it reaches the end of its input, it emits a session record for each line that still appears to be active. It normally assumes that its input is a current file, so that it uses the current time as the ending time for each session in progress. The <code>-t</code> flag causes it to use, instead, the last time found in its input, thus assuring reasonable and repeatable numbers for non-current files.
<code>-l file</code>	This option causes a line usage summary report to be placed in <i>file</i> . This report shows each line's name, number of minutes used, percentage of total elapsed time used, number of sessions charged, number of logins, and number of logins and logoffs. This report can be used to keep track of line usage, identify bad lines, and find software/hardware oddities. <i>Note that hang-up, termination of login, and termination of the login shell each generate logoff records; therefore, the number of logoffs is often three to four times the number of connect sessions.</i>
<code>-o file</code>	Using the <code>-o</code> option (for example, <code>acctcon1 -o f_overall</code> ) causes <i>file</i> to be filled with an overall record for the accounting period, giving starting time, ending time, number of reboots, and number of date changes.

The following example of the line use file (`line_use`) is created from the same `wtmp` file used in the previous `acctcon1` example; the standard output of `acctcon1` has been redirected into the file `ctmp`:

```
$ acctcon1 -t -l line_use < /etc/wtmp > ctmp
$ cat line_use
TOTAL DURATION IS 899 MINUTES
LINE      MINUTES PERCENT # SESS # ON  # OFF
console  356      95      1      1      1
tty07     69        8      1      1      1
ttya1     40        4      1      1      1
tty10     84        9      1      1      1
TOTALS   1049     --      4      4      4
```

## **prctmp**

The `prctmp` command is simple. Its only function is to put headings on the output created by `acctcon1`. `prctmp` makes a readable report from the output of `acctcon1`.

`prctmp` takes its input from standard input; therefore, to create a `prctmp` report from `acctcon1` information, you can simply pipe the output from `acctcon1` into `prctmp` as follows:

```
$ acctcon1 < /etc/wtmp | prctmp
```

`prctmp` will respond by generating a report with appropriate headings over the columns of output from `acctcon1`.

## **acctcon2**

`acctcon2` creates connect session total accounting records from standard input of the format created by `acctcon1`. In other words, to create connect session total accounting records, simply send the output from `acctcon1` into the input of `acctcon2`.

The total accounting records created by `acctcon2` are sent to standard output. So if you want to store these records, you must redirect standard output. The following command line shows how to write total accounting records from the connect session record file (`wtmp`) into the file `ctacct`:

```
$ acctcon1 < /etc/wtmp | acctcon2 > ctacct
```

---

## Process Accounting

Process accounting commands provide the means to accumulate execution statistics—such as memory usage, CPU time, number of input/output transfers—for individual processes. This section describes how to:

1. Turn process accounting on.
2. Turn process accounting off.
3. Make sure that the process accounting file (`pacct`) doesn't become too large.
4. Display process accounting records.
5. Generate a command summary report.
6. Create total accounting records from the process accounting file.

You might find it helpful to look at the System Data Flow Diagram (Figure 7-1) before reading this section.

### Turning Process Accounting On

Before System Accounting can generate process accounting data, process accounting must be turned on. Two commands can be used to accomplish this task: `turnacct on` and `accton`. After process accounting has been turned on, the kernel will write a process accounting record for every terminating process. The record will be written into the current process accounting file (`pacct` by default).

---

**Note** The `startup` command, placed in the system initialization shell script `/etc/rc`, automatically turns process accounting on. Therefore, if you have updated `/etc/rc` for System Accounting (as described in the section “How to Install System Accounting” in this chapter), process accounting will automatically be activated, and you should seldom need to use the commands described here.

These commands are described for your benefit in case you ever need to manually turn process accounting on or off.

---

## turnacct on

The command used most often to activate accounting is `turnacct on`; only the superuser and the `adm` login can execute this command. `turnacct on` assumes that the process accounting file is the default file `pacct`. The action of `turnacct on` can be summarized as follows:

1. Check to see if the process accounting file `pacct` exists.
2. If `pacct` doesn't exist, then create a new `pacct` file.
3. Turn process accounting on by invoking `accton` with the filename argument `pacct`.

To execute this command, simply enter `turnacct on` at the HP-UX prompt.

## accton

Again, only the superuser and the `adm` login can execute `accton`. When invoked with a filename argument, `accton` turns on process accounting and makes the specified filename the current process accounting file. For example,

```
$ accton /usr/adm/pacct
```

tells the kernel to start writing process accounting records to the file called `/usr/adm/pacct`. The next example would activate process accounting and make the current process accounting file `/usr/adm/XX107`:

```
$ accton /usr/adm/XX107
```

---

**Note**            The filename you specify must be an existing file; otherwise, `accton` will fail.

---

Note that in the System Data Flow Diagram (Figure 7-1), `accton` is shown calling another routine, `acct`. `acct` is the system call that actually tells the kernel to start writing process accounting records. Refer to the *HP-UX Reference* for more details on *acct(2)*.

## Turning Process Accounting Off

Two commands are used to turn process accounting off: `turnacct off` and `accton` (with no filename argument). These commands tell the kernel to stop writing records to the current process accounting file.

---

**Note** If you have updated the `/etc/shutdown` shell script as described in the section “How to Install System Accounting” in this chapter, you will seldom, if ever, use these commands. The reason is that the `shutacct` command, added to `/etc/shutdown`, automatically turns process accounting off.

---

### **turnacct off**

`turnacct off` can be executed only by the superuser and the `adm` login. `turnacct off` turns process accounting off by invoking the `accton` command without the optional filename argument. You execute this command by typing:

```
$ turnacct off
```

### **accton**

When `accton` is invoked without the optional filename argument, process accounting is turned off. You would enter this command as:

```
$ accton
```

As shown in the System Data Flow Diagram (Figure 7-1), `accton` tells the kernel to stop writing process accounting records by using the system call `acct`.

## Checking the Size of `pacct`

On a multi-user system, many processes can execute during a single hour. Therefore, process accounting files have the potential to become quite large. System Accounting has built-in mechanisms that insure that the default process accounting file `pacct` doesn't become too large. The two commands used for this purpose are: `turnacct switch` and `ckpacct`.

---

**Note**            The commands described here work only on the default process accounting file, `pacct`.

---

## **ckpacct**

The command `ckpacct` is normally invoked by `cron` every hour to insure that the current process accounting file `pacct` hasn't become too large. The format of `ckpacct` is:

```
ckpacct [blocks]
```

If the size of `pacct` exceeds the *blocks* argument, 1 000 by default if *blocks* is not specified, then `turnacct switch` is executed. `turnacct switch` renames the current `pacct` file and creates a new `pacct` file.

---

**Note**            If the amount of free space falls below a certain threshold, `ckpacct` will automatically turn off process accounting via `turnacct off`. With 5.0 HP-UX, accounting was turned off when free space fell below 2%, and back on when it went above 4%. Beginning with 5.1 HP-UX, these percentages are configured using the `acctresume` and `acctsuspend` system parameters (Refer to Appendix D, “System Parameters” and to the section “Configuring System Parameters” in Chapter 6). These percentages are in addition to the `minfree` attribute. When free space goes over the specified percentage, process accounting will be reactivated.

The kernel may also enforce a size limit on the size of `pacct`. This will take precedence over the limit set by `ckpacct`. Refer to `acctsh(1M)` and `acct(2)` in the *HP-UX Reference* for more details.

---

## **turnacct switch**

`turnacct switch` is used to create a new `pacct` file when the current `pacct` file is too large. The action of `turnacct switch` can be summarized as follows:

1. Process accounting is temporarily turned off.
2. The current `pacct` file is renamed to `pacct $incr$` , where  $incr$  is a number starting at  $1$  and incrementing by one for each additional `pacct` file that is created via `turnacct switch`.
3. After the old `pacct` file is renamed to `pacct $incr$` , a new, current `pacct` file is created.
4. Process accounting is restarted; the kernel starts writing records to the newly created `pacct` file.



The example below illustrates the effect of using the `turnacct` switch command. In the example, `turnacct switch` is executed from the `adm` home directory `/usr/adm`. Comment lines begin with a cross-hatch (`#`) and are included in the example only as explanatory material:

```
$ #
$ # First, list all the process accounting files
$ # (at this point, there is only one).
$ #
$ ll pacct*
-rw-rw-r--  1 adm    adm      2196 Mar 21 12:44 pacct
$ #
$ # Now execute turnacct switch, which will rename the current
$ # pacct file to pacct1 and will create a new pacct file.
$ #
$ turnacct switch
$ #
$ # Now verify this by listing all process accounting
$ # files again.
$ #
$ ll pacct*
-rw-rw-r--  1 adm    adm      72 Mar 21 12:46 pacct
-rw-rw-r--  1 adm    adm     2196 Mar 21 12:44 pacct1
$ #
$ # The current process accounting file is pacct.  The previous
$ # process accounting file is now named pacct1.
$ #
```

## Displaying Process Accounting Records – `acctcom`

The `acctcom` command allows you to display records from any file containing process accounting records. Normally you would use this command to display records from the `pacct` files (`pacct`, `pacct1`, `pacct2` ... ).

`acctcom` is a very versatile command; its syntax follows:

```
acctcom [ options ] [ file ] ...
```

If no *file* is specified, `acctcom` uses the current `pacct` file as input. Input can also be taken from standard input. Some of `acctcom`'s options allow you to select

only the records that you want to see; other options control the format of the report.

The information contained in this section is organized as follows:

- First, definitions are given for the columnar data produced by `acctcom`.
- Command options that control the format of the report are discussed.
- Options that allow you to select particular records are described.
- Finally, to help you understand how to use `acctcom`'s options, sample `acctcom` reports are shown.

### Definitions of Information Produced by `acctcom`

`acctcom` generates a columnar report with descriptive headings on each column. Each line of the report represents the execution statistics that a particular process accumulated during its lifetime. The standard columns in the report—that is, the columns that are displayed when none of `acctcom`'s options are specified—are shown below:

Column Header	Definition
COMMAND NAME	The name of the command or program that spawned the process is shown here. Whenever you enter a command, you are spawning a process. For example, if you enter the command <pre>\$ ll /usr/lib/acct</pre> you are creating a process with the command name <code>ll</code> . If a command requiring superuser privileges is executed, a <code>#</code> appears before the command name.
USER	The login name of the user who created the process is displayed here.
TTYNAME	This is the name of the terminal from which the process was executed. If the process was not executed from a known terminal (for example, if it was executed via <code>cron</code> ), then a question mark(?) appears in this column.
START TIME	The time that the process began executing (in <code>hh:mm:ss</code> format) is displayed here.

<b>Column Header</b>	<b>Definition</b>
END TIME	This is the time ( <i>hh:mm:ss</i> ) that the process finished executing.
REAL (SECS)	The number of seconds that elapsed from START TIME to END TIME is shown in this column.
CPU (SECS)	This column shows how much of the CPU's time a process used during its execution.
MEAN SIZE(K)	This is a rough estimate (in kilobytes) of the amount of memory that a process used during execution.  This estimate is determined from the current process's memory usage at each system clock interrupt. It is, therefore, subject to statistical sampling errors. Only the memory resident pages of a process are counted; no pages in the swap space are counted. Shared code and data is divided among the processes using it. The size is divided by the number of processes sharing the code or data.

Listed below are the columns that are not displayed on the standard report, but which can be displayed by using `acctcom` options:

<b>Column Header</b>	<b>Definition</b>
F	For a process created by <code>fork</code> which does not do an <code>exec</code> , this column takes the value <i>1</i> ; commands which require superuser privileges show a <i>2</i> ; superuser commands which do a <code>fork</code> without an <code>exec</code> show a <i>3</i> ; otherwise, this column shows a <i>0</i> .
STAT	This column displays the system exit status. (This is <i>not</i> the status returned by <code>exit</code> to a parent process during <code>wait</code> ). When a process terminates normally, this field shows a <i>0</i> . If a command terminates abnormally, then a value other than zero is shown. For example, if you interrupt a command with the <code>DEL</code> key, this column will contain a <i>2</i> .
HOG FACTOR	The hog factor is computed as the CPU time divided by REAL time; it provides a relative measure of the available CPU time used by the process during its execution. For example, a hog factor of less than 0.50 indicates that the process spent less than half of its time using the CPU. A hog factor of 0.75 indicates that a process spent 75% of its time using the CPU.

<b>Column Header</b>	<b>Definition</b>
KCORE MIN	<p>This calculation provides a combined measurement of the amount of memory used (in kilobytes) and the length of time it was used (in minutes). It is computed as follows:</p> $\text{KCORE MIN} = \text{CPU (SECS)} * \text{MEAN SIZE(K)} / 60$
CPU SYS	<p>This is the portion of total CPU time that was spent executing operating system code, such as system calls (for example, writing to disk).</p>
USER (SECS)	<p>This is the remaining portion of CPU time. User CPU time is the amount of time actually spent executing a process's code (rather than system code).</p>
CPU FACTOR	<p>Whenever you execute a command, the CPU spends part of its time actually executing the command's code (user CPU time) and spends the rest of its time performing system functions, such as writing to the disk or terminal (system CPU time). That is, total CPU time is comprised of both CPU SYS and USER CPU time:</p> $\text{CPU (SECS)} = \text{CPU SYS} + \text{USER (SECS)}$ <p>The CPU factor shows the ratio of user CPU time to total CPU time:</p> $\text{CPU FACTOR} = \text{USER (SECS)} / (\text{CPU SYS} + \text{USER (SECS)})$ <p>For example, if a command has a CPU factor of 0.35, that means that the CPU spent 35% of its time executing user code and 65% performing system functions.</p>
CHARS TRNSFD	<p>The number of characters (bytes) read and/or written by the command is displayed in this column.</p>
BLOCKS R/W	<p>This column shows the number of file system blocks read and/or written as a result of executing this command. <i>This number is not directly related to CHARS TRNSFD and may vary each time the command is executed</i>, because BLOCKS R/W is affected by directory searches made before opening files, other processes accessing the same files, and general file system activity.</p>

## Report Format Options

When no report format options are specified, `acctcom` will produce a report containing only the default information. Optional information can be displayed only by using the report format options. Definitions of the report format options follow:

Option	Description
-a	Cause average statistics to be displayed at the end of the report. The following information is shown: total number of commands processed ( <code>cmds=xxx</code> ) <ul style="list-style-type: none"><li>■ average real time per process (<code>Real=x.xx</code>)</li><li>■ average CPU time per process (<code>CPU=x.xx</code>)</li><li>■ average USER CPU time per process (<code>USER=x.xx</code>)</li><li>■ average SYS CPU time per process (<code>SYS=x.xx</code>)</li><li>■ average characters transferred (<code>CHARS=x.xx</code>)</li><li>■ average blocks transferred (<code>BLK=x.xx</code>)</li><li>■ average CPU factor (<code>USR/TOT=x.xx</code>)</li><li>■ average HOG factor (<code>HOG=x.xx</code>)</li></ul>
-b	Display the process records in reverse order: most recently executed commands will be shown first.
-f	Print the <code>fork/exec</code> flag (F column) and process exit status (STAT column) on the report.
-h	Cause the optional HOG FACTOR column to be displayed, instead of the standard mean memory size column MEAN SIZE(K).
-i	Replace the standard MEAN SIZE(K) column in the report with the optional I/O counts—CHARS TRNSFD and BLOCKS R/W.
-k	Replace the standard MEAN SIZE(K) column with KCORE MIN.
-m	Show the default column MEAN SIZE(K) on the report. This option is used to include MEAN SIZE(K) when it has been bumped off by another option. For example: <pre>\$ acctcom -km</pre> produces a report showing both KCORE MIN and MEAN SIZE(K).

Option	Description
-r	Include the optional CPU FACTOR column in the report.
-t	Show separate system and user CPU times (CPU SYS and USER (SECS), respectively).
-v	Suppress the printing of column headings at the top of the report.
-q	This option is the same as the -a option, except that individual process accounting records are not displayed—only the averages are displayed.
-o <i>ofile</i>	Copy the input process accounting records to <i>ofile</i> .

## Record Selection Options

The options described here allow you to select the records that are included in the report produced by `acctcom`. For each option, descriptions and examples are provided:

Option	Description
-l <i>line</i>	<p>Display only the processes that were executed from the user terminal <i>/dev/line</i>. For example:</p> <pre>\$ acctcom console</pre> <p>would display records only for the processes that were created from the terminal <code>console</code>.</p>
-u <i>user</i>	<p>Show only the processes belonging to <i>user</i>. <i>user</i> can be any of the following:</p> <ul style="list-style-type: none"> <li>■ a user ID (for example, <code>acctcom -u 355</code>)</li> <li>■ user name (<code>acctcom -u julie</code>)</li> <li>■ a cross-hatch <code>#</code> (<code>acctcom -u#</code>)</li> <li>■ a question mark <code>?</code> (<code>acctcom -u?</code>)</li> </ul> <p>If <code>#</code> is specified as the user name, then only the commands that require superuser privileges will be displayed by <code>acctcom</code>. If <code>?</code> is given as the user, then only the processes with unknown process IDs will be displayed. As an example, the following two commands are equivalent:</p> <pre>\$ acctcom -u 0 \$ acctcom -u root</pre>

<b>Option</b>	<b>Description</b>
<b>-g</b> <i>group</i>	<p>Show only the processes belonging to <i>group</i>. <i>group</i> may be specified as either a group name or group ID. For example, suppose the group <i>pseudo</i> with group ID 300 is defined in <i>/etc/group</i>; then the following two commands are equivalent:</p> <pre style="margin-left: 40px;">\$ acctcom -g 300 \$ acctcom -g pseudo</pre>
<b>-s</b> <i>time</i>	<p>Select processes existing <i>at or after time</i>. Time is given in 24-hour format—<i>hr[:min[:sec]]</i>. The following example would display all the processes that existed at or after 3:30pm:</p> <pre style="margin-left: 40px;">\$ acctcom -s 15:30</pre>
<b>-e</b> <i>time</i>	<p>Select processes that existed <i>at or before time</i>. Time is supplied in 24-hour format <i>hr[:min[:sec]]</i>. The next example would display all the processes that existed between midnight and 12:15am:</p> <pre style="margin-left: 40px;">\$ acctcom -e 0:15</pre>
<b>-S</b> <i>time</i>	<p>Select processes <i>starting</i> at or after time where <i>time</i> is in 24-hour format. The following example would display all the processes that <i>started</i> at 1:30:42pm or after:</p> <pre style="margin-left: 40px;">\$ acctcom -S 13:30:42</pre>
<b>-E</b> <i>time</i>	<p>Display only the processes that <i>terminated</i> at or before <i>time</i>, where time is in 24-hour format <i>hr[:min[:sec]]</i>. Note both the <b>-S</b> and <b>-E</b> options with the same <i>time</i> argument will cause <i>acctcom</i> to display only the processes that existed at the specified <i>time</i>. For example, to see all the processes that existed at exactly 30 minutes past noon, you would enter:</p> <pre style="margin-left: 40px;">\$ acctcom -S 12:30 -E 12:30</pre>

Option	Description
-n <i>pattern</i>	<p>Show only the commands matching <i>pattern</i>. <i>pattern</i> can be a regular expression as described in <i>ed(1)</i>, except that <i>+</i> means one or more occurrences. For example, to display all processes that were created by executing the <code>ls</code> command, you would enter:</p> <pre style="margin-left: 40px;">\$ acctcom -n ls</pre> <p>To display all the commands that start with <code>acct</code>, enter:</p> <pre style="margin-left: 40px;">\$ acctcom -n acct</pre> <p>To see all the commands that contain the letter <i>m</i> in their spelling you can use the wildcard character <code>*</code>. Type:</p> <pre style="margin-left: 40px;">\$ acctcom -n .*m.*</pre>
-H <i>factor</i>	<p>Display only those processes whose hog factor exceeds <i>factor</i>. For example,</p> <pre style="margin-left: 40px;">\$ acctcom -H 0.85</pre> <p>would display all the processes that spent over 85% of their execution time in the CPU. You can use this option to find greedy processes—processes that are hogging the CPU.</p>
-O <i>time</i>	<p>Show only those processes whose system CPU time exceeds <i>time</i>, specified in seconds. The following example would be used to determine which processes took more than 8.25 seconds of operating system CPU time to execute:</p> <pre style="margin-left: 40px;">\$ acctcom -O 8.25</pre> <p>This option could be used to determine which processes are making heavy use of the operating system calls.</p>
-C <i>sec</i>	<p>Show only the processes whose total CPU time (<code>SYS + USER</code>) exceeds <i>sec</i> seconds. The next example would display all the processes that used over 5.28 seconds of CPU time to execute:</p> <pre style="margin-left: 40px;">\$ acctcom -C 5.28</pre>
-I <i>chars</i>	<p>Display only the processes transferring more characters than the limit given by <i>chars</i>. For example,</p> <pre style="margin-left: 40px;">\$ acctcom -I 10240</pre> <p>will display all the processes that transferred over ten kilobytes of characters (<math>10\ 240 = 10 \times 1\ 024</math> bytes).</p>



## Sample Reports

The following sample report illustrates the use of `acctcom` without any options. The report generated is the standard report produced when no options are specified:

```
$ acctcom
ACCOUNTING RECORDS FROM: Thu Mar 21 12:52:26 1985
COMMAND          START      END          REAL      CPU      MEAN
NAME            USER      TTYNAME     TIME      TIME      (SECS)     (SECS)    SIZE(K)
#accton        root      console    12:52:26 12:52:26    0.12      0.10     19.00
ls             sarah     tty07      14:04:08 14:04:08    0.28      0.23     16.50
ckpacct       adm       ?          14:30:00 14:30:05    5.13      1.45     24.00
pwd           bill      tty10      15:09:07 15:09:07    0.48      0.22     22.50
find          sarah     tty07      18:51:37 18:51:39    2.73      0.15     26.50
tabs          root      console    19:10:18 19:10:18    0.92      0.13     23.50
stty          root      console    19:10:19 19:10:19    0.88      0.08     26.00
mail          bill      tty10      19:10:21 19:10:22    1.78      0.23     28.50
news          root      console    19:10:23 19:10:23    0.73      0.12     23.00
acctcom       adm       ttya0      19:53:16 19:53:38    22.58     2.55     28.50
```

The next example displays all the processes created between 7:00pm and 7:30pm by the user `root`. In addition, the optional CPU factor and average statistics are included in the output:

```
$ acctcom -S 19:00 -E 19:30 -u root -ah
START AFT: Thu Mar 21 19:00:00 1985
END BEFOR: Thu Mar 21 19:30:00 1985
COMMAND          START      END          REAL      CPU      HOG
NAME            USER      TTYNAME     TIME      TIME      (SECS)     (SECS)    FACTOR
tabs          root      console    19:10:18 19:10:18    0.92      0.13     0.14
stty          root      console    19:10:19 19:10:19    0.88      0.08     0.09
news          root      console    19:10:23 19:10:23    0.73      0.12     0.16
cmds=3 Real=0.84 CPU=0.11 USER=0.02 SYS=0.09 CHAR=26.12 BLK=11.50
USR/TOT=0.19 HOG=0.13
```

Sample reports are helpful, but the best way to learn the various `acctcom` options is to use them. Take a few minutes to experiment with this command; it is very powerful and can provide you with much useful information if used properly.

## Command Summary Report – `acctcms`

The `acctcms` command takes process accounting records as input; but instead of reporting on the individual processes, `acctcms` generates a report on the commands that generated the process accounting records. The action of `acctcms` can be summarized as follows:

1. `acctcms` looks through the input process accounting records and accumulates execution statistics for each unique command name. This information is stored in internal summary format—one record per command name.
2. Depending on the `acctcms` options used, the command summary records created in step 1 are sorted.
3. The command summary records are written to standard output in the internal summary format mentioned in step 1. This format is not readable.

The syntax of the `acctcms` command is:

```
acctcms [options] files
```

where *files* is a list of the input process accounting files for which the command summary report is to be generated. The *options* are discussed in the following sections.

### Producing a Readable Report – the `-a` option

By default, the output of `acctcms` is in internal summary record format; if you display it to your terminal, all you see is gibberish. To get a ASCII, readable report, use the `-a` option.

The `-a` option causes `acctcms` to produce a report with descriptive column headings. Total and average (mean) execution statistics for each command are displayed—one line per command—along with total and average statistics over all commands in the report. Descriptions of the columnar data produced by `acctcms` are shown in the following table.

<b>Column Header</b>	<b>Description</b>
COMMAND NAME	The name of the command for which execution statistics are summarized. Unfortunately, all shell procedures are lumped together under the name <code>sh</code> , because only object modules are reported by the process accounting system.
NUMBER CMDS	The total number of times that the command was invoked.
TOTAL KCOREMIN	The total amount of kcore minutes accumulated for the command. (Refer to the section “Definitions of Information Produced by acctcom” in this chapter for a more complete description of kcore minutes.)
TOTAL CPU-MIN	The total CPU time that the named command has accumulated.
TOTAL REAL-MIN	Total accumulated real time seconds are displayed in this column.
MEAN SIZE-K	The average amount of memory (in kilobytes) consumed by the command.
MEAN CPU-MIN	The average CPU time consumed per command invocation is shown here; the following equation shows how it is computed: $\text{MEAN CPU-MIN} = \text{TOTAL CPU-MIN} / \text{NUMBER CMDS}$
HOG FACTOR	The average hog factor over all invocations of the command. It is computed as: $\text{HOG FACTOR} = \text{TOTAL CPU-MIN} / \text{TOTAL REAL-MIN}$
CHARS TRNSFD	The total number of characters transferred by the command. Note that this number may sometimes be negative.
BLOCKS READ	A total count of the physical blocks read and written by the given command. (Refer to the section “Displaying Process Accounting Records—acctcom” in this chapter for details on the significance of this total.)

---

**Note**

When only the `-a` option is specified, the report is sorted in descending order on the `TOTAL KCOREMIN` column: commands using more `TOTAL KCOREMIN` are shown before those using fewer `TOTAL KCOREMIN`. This report gives a relative measure of the amount of memory used over time by the various commands: commands toward the start of the report are making more use of memory resources than are commands toward the end of the report.

---

**Other Options**

In addition to the `-a` option, several other options can be used to control the format of the report generated by `acctcms`. Some options specify which field to sort the report on; other options control the printing of prime/non-prime time usage. The options and a description of their use follow:

<b>Option</b>	<b>Description</b>
<code>-c</code>	Sort the commands in descending order on <code>TOTAL CPU-MIN</code> , rather than the default <code>TOTAL KCOREMIN</code> . This report can be used to determine which commands are using most of the computer's CPU time.
<code>-n</code>	Cause the report to be sorted in descending order on the column named <code>NUMBER CMDS</code> . Commands toward the start of this report are the ones used most frequently; commands toward the end are used least often.
<code>-j</code>	Combine all commands invoked only once on one line of the report. This line is denoted by having <code>***other</code> in the <code>COMMAND NAME</code> column. This option is useful for shortening a report that has many one-invocation commands.
<code>-o</code>	<i>Used only with the <code>-a</code> option</i> , <code>-o</code> causes the report to be generated only for commands that were executed during non-prime time (as specified in the <code>holidays</code> file). You can use this option to get a non-prime time command summary report.
<code>-p</code>	<i>Also used only with the <code>-a</code> option</i> , <code>-p</code> elicits a report generated only for commands that were executed during prime time (as specified in <code>holidays</code> ). This option is used to get a prime time command summary report.

<b>Option</b>	<b>Description</b>
-apo	When the options -o and -p are used together with -a, a combination prime and non-prime time report is produced. The output of this report is same as that produced by -a alone, except that the NUMBER CMDS, TOTAL CPU-MIN, and TOTAL REAL-MIN columns are divided into two columns—one for prime time totals, the other for non-prime time. (Prime time columns have a (P) header, while non-prime time columns are headed by (NP).)
-s [ <i>files</i> ]	Specifies that any named input <i>files</i> following the -s on the command line are already in internal summary format. This option is useful for merging previous acctcms reports with current reports. The following example uses -s to create a command summary report from previous process accounting files (pacct?) and the current process accounting file (pacct). The final ASCII report is stored in the file ascii_summary. <pre>\$ acctcms pacct? &gt; old_summary \$ acctcms pacct &gt; new_summary \$ acctcms -as old_summary new_summary &gt; ascii_summary</pre>

## Sample Report

The ASCII reports produced by acctcms contain more than 80 characters per line. When these reports are displayed at an 80-column terminal, the lines wrap around on the screen. In addition, if the report is printed on an 80-column printer, some of the rightmost columns will be lost. Therefore, be sure to do one of the following:

- Use a printer with compressed print capabilities, so that all of the report will fit on standard computer paper.
- Use a printer with enough columns to display all the information—for example, a 132-column printer.

The following example generates a command summary report for the current process accounting file (no file is specified, so the current `pacct` file is assumed). By giving the `-j` option, all the commands that were executed only once are grouped under the command name `***other`. Note also that total execution statistics for all commands are grouped under the command name `TOTALS`.

```
$ acctcms -aj
```

TOTAL COMMAND SUMMARY									
COMMAND NAME	NUMBER CMDS	TOTAL KCOREMIN	TOTAL CPU-MIN	TOTAL REAL-MIN	MEAN SIZE-K	MEAN CPU-MIN	HOG FACTOR	CHARS TRNSFD	BLOCKS READ
TOTALS	61	17.63	0.38	164.49	46.25	0.01	0.00	104553	1027
acctcms	17	12.13	0.16	0.35	76.72	0.01	0.45	49192	306
sh	8	2.43	0.09	152.86	26.79	0.01	0.00	9043	163
more	3	0.73	0.02	10.50	31.00	0.01	0.00	21618	83
ll	6	0.61	0.04	0.11	16.50	0.01	0.33	5715	95
acctcom	4	0.58	0.02	0.07	28.50	0.01	0.30	15319	42
***other	9	0.54	0.02	0.14	25.26	0.00	0.16	459	161
cat	4	0.19	0.01	0.35	22.97	0.00	0.02	3112	52
rm	2	0.11	0.00	0.02	22.22	0.00	0.29	0	29
chmod	2	0.10	0.00	0.01	22.00	0.00	0.35	0	15
accton	2	0.08	0.00	0.02	19.00	0.00	0.29	0	22
sed	2	0.08	0.01	0.04	14.50	0.00	0.13	73	38
echo	2	0.05	0.00	0.02	20.00	0.00	0.16	22	21

## Creating Total Accounting Records

Two commands—`acctprc1` and `acctprc2`—are used to create total accounting records from the process accounting files. The output from `acctprc1` is supplied as input to `acctprc2` which produces the total accounting records. These commands are normally invoked by `runacct` to produce daily accounting information.

## acctprc1

This command reads process accounting records from standard input, adds login names corresponding to the user ID of each record, and then writes for each process an ASCII line showing:

- the ID of the user that created the process
- the user's login name
- prime CPU time in ticks (a “tick” is one fiftieth of a second)
- non-prime CPU time, also in ticks
- mean memory size (in pages—4 Kbytes on Series 300)

The format of `acctprc1` is:

```
acctprc1 [ctmp]
```

where `ctmp` contains a list of login sessions of the form created by `acctcon1`, sorted by user ID and login name.

To use `acctprc1`, input must be redirected from a process accounting file. The following example creates a file, `ascii_ptacct`, containing ASCII process accounting information that can be used to create process total accounting records. This file is created from the current process accounting file `pacct`.

```
$ acctprc1 <pacct >ascii_ptacct
```

Normally, `acctprc1` gets login names from the password file `/etc/passwd`, which is sufficient on systems where each user has a unique user ID. However, on systems where different users share the same user ID, the `ctmp` file should be specified; it helps `acctprc1` distinguish different login names that share the same user ID.

## acctprc2

This command reads from standard input records of the form created by `acctprc1`; it then summarizes the records by user ID and name, and writes the sorted summaries to standard output as total accounting records. The following example creates total accounting records for all processes in the current process accounting file `pacct`; the total accounting records are stored in the file `ptacct`.

```
$ acctprc1 <pacct |acctprc2 >ptacct
```

---

## Charging Fees to Users – chargefee

System Accounting provides the capability to charge fees to specific users; the `chargefee` command is used to accomplish this task. `chargefee` allows you to charge generic *units* to a specific login name. The syntax of this command is:

```
chargefee login_name number
```

where *number* is the number of units to be charged to a particular user, and *login\_name* is the login name of the user to whom *number* units are to be charged.

---

**Note**            *number* can be any whole number in the range -32 768 to 32 767; when charging fees, keep in mind that the sum of each user's fees must also be within this range.

---

`chargefee` accumulates fee charge records in the file `/usr/adm/fee`. These records are then merged with other accounting records via `runacct`.

The following example charges 25 units to the user whose login name is `julie`:

```
$ chargefee julie 25
```

Suppose you inadvertently charged 247 units to the user named `zimblits`, and you want to return his charges to their original value. You would enter the following:

```
$ chargefee zimblits -247
```



---

## Summarizing and Reporting Accounting Information

This final group of commands summarizes and reports accounting information. Certain commands display and merge total accounting files, while others generate the daily and monthly reports used to analyze system performance and bill users for resource usage. The following commands are discussed in this section:

- `prtacct`—displays total accounting records
- `acctmerg`—merges total accounting files
- `runacct`—generates daily summary files and reports
- `prdaily`—displays the daily summary files and reports created by `runacct`
- `monacct`—creates monthly summary files and reports

### Displaying Total Accounting Records – `prtacct`

The `prtacct` command allows you to display the contents of a process accounting file. Its format is

```
prtacct file "heading"
```

where:

- *file* is the name of the total accounting file to be displayed.
- “*heading*” is a comment to be included in the standard report header produced by `prtacct`.

The format of the `prtacct` report is described next and is followed by an example.

## Report Format

prtacct produces a columnar report with one line per total accounting record. Descriptive column headings are included in the report. Definitions of each column follow:

<b>Column Header</b>	<b>Description</b>
UID	This column shows the user ID of the owner of the total accounting record—that is, the ID of the user for whom the total accounting record was created.
LOGIN NAME	The login name of the owner of the total accounting record is displayed here.
CPU (MINS)	This column shows the total amount of CPU time (in minutes) that the user has consumed. This column is divided into prime and non-prime columns (PRIME and NPRIME, respectively). Information in these columns is created through process accounting commands.
KCORE-MINS	This represents a cumulative measure of memory and CPU time that a user consumed (refer to the section “Definitions of Information Produced by acctcom” in this chapter for a more precise definition). Information in this column is also divided into PRIME and NPRIME columns. This information is created through process accounting commands.
CONNECT (MINS)	This identifies the real time used (in minutes). In essence, what this column identifies is the amount of time that the user was logged in to the system. This column is also subdivided into PRIME and NPRIME columns. The connect session accounting commands are the source of this information.
DISK BLOCKS	The total number of disk blocks allocated to the user is shown here. This information is created via disk space accounting commands.
# OF PROCS	The total number of processes spawned by the user is displayed here. This information is created via the process accounting commands.

<b>Column Header</b>	<b>Description</b>
# OF SESS	This column shows how many times the user logged in. Connect session accounting commands create this data.
# DISK SAMPLES	This column indicates how many times the disk accounting was run to obtain the average number of disk blocks listed in the DISK BLOCKS column.
FEE	The number of fee units charged via <code>chargefee</code> is displayed here.

### **Example**

The following example displays disk total accounting records. First, the total accounting records are created via disk space accounting commands; then, they are displayed using `prtacct`. When examining this report, take note of the following:

- There are many similarities between this and the sample report produced by `diskusg` (refer to the section “Disk Space Usage Accounting” in this chapter).
- Only the columns relating to disk space usage have non-zero values, because the total accounting records were created only from disk space usage accounting commands.

The example report produced by prtacct follows:

```
$ for file_system in `cat /etc/checklist`
> do
> diskusg $file_system >dtmp.`basename $file_system`
> done
$ diskusg -s dtmp.* |sort +0n +1 |acctdisk >disktacct
$ prtacct disktacct "DISC TOTAL ACCOUNTING RECORDS"
```

Mar 26 17:01 1985 DISC TOTAL ACCOUNTING RECORDS Page 1

UID	LOGIN	CPU (MINS)		KCORE-MINS		CONNECT (MINS)		DISK	# OF	# OF	#	DISK	FEE
	NAME	PRIME	NPRIME	PRIME	NPRIME	PRIME	NPRIME	BLOCKS	PROCS	SESS	SAMPLES		
0	TOTAL	0	0	0	0	0	0	11598	0	0	10	0	0
0	root	0	0	0	0	0	0	10616	0	0	1	0	0
1	bin	0	0	0	0	0	0	778	0	0	1	0	0
4	adm	0	0	0	0	0	0	96	0	0	1	0	0
350	fred	0	0	0	0	0	0	14	0	0	1	0	0
351	bill	0	0	0	0	0	0	32	0	0	1	0	0
352	mike	0	0	0	0	0	0	20	0	0	1	0	0
353	sarah	0	0	0	0	0	0	16	0	0	1	0	0
354	molly	0	0	0	0	0	0	22	0	0	1	0	0
355	julie	0	0	0	0	0	0	2	0	0	1	0	0
501	guest	0	0	0	0	0	0	2	0	0	1	0	0

## Merging Total Accounting Files – acctmerg

Normally executed by runacct, the acctmerg command merges separate total accounting files into a single total accounting file. All the total accounting records for a particular user name and ID are merged together to form one total accounting record for the given user name and ID. This command is useful for merging disk, connect session, and process total accounting files together to form a single, comprehensive total accounting file.

`acctmerg` reads standard input and up to nine additional files, all in total accounting record format. Its syntax is:

```
acctmerg [options] [file] ...
```

where:

- *options* control the report format and the manner in which records are merged.
- *file* is one of up to nine files (in addition to standard input) that are to be merged into a single total accounting file, written to standard output.

## Command Options

The following options may be used with `acctmerg` to control the report format and the manner in which the total accounting records are merged:

Option	Description
-a	<code>acctmerg</code> normally produces output as total accounting records. The <code>-a</code> option causes <code>acctmerg</code> to produce output in ASCII. Note that the output generated by using this option is the same as the report produced by <code>prtacct</code> , except that no report headings or totals are displayed—only the columnar data is shown.
-i	In the default case, <code>acctmerg</code> assumes that its input files contain total accounting records. If <code>-i</code> is specified, then <code>acctmerg</code> will expect input files to be in the ASCII format created by the <code>-a</code> option.
-p	This option simply echoes input records—no merging or processing is done. The output is displayed in the format produced by the <code>-a</code> option.
-t	This option produces a single total accounting record that summarizes all input records. To see the ASCII version of this record, you must use the <code>-t</code> and <code>-a</code> options together: <pre>\$ acctmerg -t -a &lt;tot_acct_recs</pre> <i>Note that <code>-t</code> and <code>-a</code> can be specified in any order, but they must be specified separately as shown.</i>
-u	Normally, <code>acctmerg</code> merges records that have the same user ID and user name. Using <code>-u</code> causes <code>acctmerg</code> to merge records on the basis of same user ID only—that is, the user name is disregarded as a key on which to merge records.
-v	This option causes <code>acctmerg</code> to produce output in verbose ASCII format. The same report is produced as the <code>-a</code> option, except that floating point numbers are displayed in more precise notation: <i>&lt;mantissa&gt;e&lt;exponent&gt;</i>

The `-a`, `-v`, and `-i` options are useful if you wish to edit total accounting records. For example, suppose that you have created a total accounting file (`ptacct`) containing process total accounting records, and you want to make

some adjustments to these records. The following sequence could be used to “repair” to this file:

```
$ acctmerg -v -a <ptacct >ptacct.ascii
    edit ptacct.ascii as desired ...
    then copy the changes back to ptacct
$ acctmerg -i <ptacct.ascii >ptacct
```

## Example

The following example creates disk, process, and connect session total accounting records, merges them together, and stores the merged file in the file merged\_file:

```
$ # First, create disk space usage total accounting records (dtacct) ...
$ #
$ for fs in `cat /etc/checklist`
> do
> diskusg $fs >dtmp.`basename $fs`
> done
$ diskusg -s dtmp.* |sort +0n +1 |acctdisk >dtacct
$ #
$ # Now create connect session total accounting records (ctacct) ...
$ #
$ acctcon1 </etc/wtmp |acctcon2 >ctacct
$ #
$ # Create process total accounting records (ptacct) ...
$ #
ptacct
$ for p_file in pacct*
> do
> acctprc1 <$p_file |acctprc2 >>ptacct
> done
$ #
$ # Now merge all the total accounting files (?tacct) into
$ # a single total accounting file (tacct) ...
$ #
$ acctmerg dtacct ctacct <ptacct >tacct
```

## Creating Daily Accounting Information – runacct

`runacct` is the main daily accounting shell procedure. Start `runacct` via `cron` during non-prime hours, when users are logged off. This is because it does not correctly log time for users that log on before running `runacct`.

`runacct` processes disk, connect session, process, and fee accounting files. It prepares cumulative summary files for use by `prdaily` and for billing purposes. This section discusses the following aspects of `runacct`:

- files processed by `runacct`
- the states that `runacct` progresses through while executing
- recovery from `runacct` failure
- restarting `runacct`
- reports produced by `runacct`

### Files Processed by runacct

The following files, processed by `runacct`, are of particular interest to the reader. (Filenames are given relative to the directory `/usr/adm/acct`.)

- `nite/lineuse` contains usage statistics for each terminal line on the system. This report is especially useful for detecting bad lines. If the ratio of logoffs to logins on a particular line exceeds 3 to 1, then there is a good possibility that the line is failing.
- `nite/daytacct` contains total accounting records from the previous day.
- `sum/tacct` contains accumulated total accounting records for each day's total accounting records (`nite/daytacct`) and can be used for billing purposes. It is restarted each month or fiscal period by the `monacct` shell script.
- `sum/daycms` is produced by `acctcms`. It contains the daily command summary. The ASCII version of this file is in `nite/daycms`.
- `sum/cms` holds the accumulation of each day's command summaries (`sum/daycms`). A new `sum/cms` file is created each month by `monacct`. The ASCII version of this file is in `nite/cms`.
- `sum/loginlog` maintains a record of the last time each user logged in.
- `sum/rprtmmdd` is the main daily accounting report created by `runacct`. The name for this report is created automatically by the system with `mm`



being the month and *dd* the day of the report. This report can be printed via `prdaily`.

`runacct` takes care not to damage files in the event of errors. A series of protection mechanisms are used that attempt to recognize errors, provide intelligent diagnostics, and terminate processing in such a way that `runacct` can be restarted with minimal intervention. To accomplish these goals, the following actions are performed by `runacct`:

- `runacct`'s progress is recorded by writing descriptive messages to the `nite/active` file.
- All diagnostics output during the execution of `runacct` are redirected to the file `nite/fd2log`.
- If the files `lock` and `lock1` exist when `runacct` is invoked, an error message will be displayed, and execution will terminate.
- The `lastdate` file contains the month and day that `runacct` was last run and is used to prevent more than one execution per day.
- If `runacct` detects an error, a message is written to `/dev/console`, mail is sent to `root` and `adm`, locks are removed, diagnostics files are saved, and execution is terminated.

## The States of `runacct`

In order to allow `runacct` to be restartable, processing is broken down into separate re-entrant *states*. As `runacct` executes, it records its progress by writing the name of the most recently completed state into the file called `/usr/adm/statefile`. After processing for a state is complete, `runacct` examines `statefile` to determine which state to enter next. When `runacct` reaches the final state (`CLEANUP`), the `lock` and `lock1` files are removed, and execution terminates. `runacct`'s states are described in the next table.

State	Action
SETUP	The command <code>turnacct</code> switch is executed. The process accounting files, <code>pacct?</code> , are moved to <code>Spacct?.mddd</code> . The <code>/etc/wtmp</code> file is moved to <code>nite/wtmp.mddd</code> with the current time added on the end.
WTMPFIX	<code>nite/wtmp.mddd</code> is checked for correctness by <code>wtmpfix</code> . Some date changes will cause <code>acctcon1</code> to fail, so <code>wtmpfix</code> attempts to adjust the time stamps in the <code>nite/wtmp.mddd</code> file if a date change record appears.
CONNECT1	Connect session records are written to <code>ctmp</code> . The <code>lineuse</code> file is created, and the <code>reboots</code> file, showing all of the boot records found in <code>nite/wtmp.mddd</code> , is created.
CONNECT2	<code>ctmp</code> is converted to connect session total accounting records in the file <code>ctacct.mddd</code> .
PROCESS	The <code>acctprc1</code> and <code>acctprc2</code> programs are used to convert the process accounting files, <code>Spacct?.mddd</code> , to the total accounting records in <code>ptacct?.mddd</code> . The <code>Spacct</code> and <code>ptacct</code> files are correlated by number so that if <code>runacct</code> fails, the unnecessary reprocessing of <code>Spacct</code> files will not occur. One precaution should be noted: <i>when restarting runacct in this state, remove the last ptacct file; if you don't, runacct will not finish.</i>
MERGE	Merge the process and connect session total accounting records to form <code>nite/daytacct</code> .
FEES	Merge in any ASCII <code>tacct</code> records from the file <code>fee</code> into <code>nite/daytacct</code> .
DISK	On the day after the <code>dodisk</code> shell script runs, merge <code>nite/disktacct</code> with <code>nite/daytacct</code> .
MERGETACCT	Merge <code>nite/daytacct</code> with <code>sum/tacct</code> , the cumulative total accounting file. Each day, <code>nite/daytacct</code> is saved in <code>sum/tacctmddd</code> , so that <code>sum/tacct</code> can be recreated in the event it becomes corrupted or lost.
CMS	Merge in today's command summary with the cumulative summary file <code>sum/cms</code> . Produce ASCII and internal format command summary files.

<b>State</b>	<b>Action</b>
USEREXIT	Any installation-dependent (local) accounting programs can be run in this state. For example, you might want to execute commands that generate daily billing data for individual users (the shell script <code>acct_bill</code> in the section “Sample System Accounting Shell Scripts” could be used for this purpose). To have local accounting programs executed by <code>runacct</code> , simply enter the commands in <code>runacct</code> in the code for the USEREXIT state of <code>runacct</code> .
CLEANUP	Clean up the temporary files, run <code>prdaily</code> and save its output in the file <code>sum/rprtmmdd</code> , remove the locks, then exit.

## Recovering from Failure

It is possible that `runacct` might fail and terminate abnormally. The primary reasons for `runacct` failure are:

- a system “crash”
- not enough disk space remaining in `/usr`
- a corrupted `wtmp` file

If the `nite/activemmdd` file exists, check it first for error messages. If the `nite/active` file and lock files exist, check `fd2log` for any mysterious messages. The following are error messages produced by `runacct` and the recommended recovery actions:

```
ERROR: locks found, run aborted
```

The files `lock` and `lock1` were found. These files must be removed before `runacct` can be restarted.

```
ERROR: acctg already run for date: check /usr/adm/acct/nite/lastdate
```

The date in `lastdate` and today’s date are the same. Remove `lastdate` before restarting `runacct`.

```
ERROR: turnacct switch returned rc=?
```

Check the integrity of `turnacct` and `accton`. The `accton` program must be owned by `root` and have the `setuid` bit set.

```
ERROR: Spacct?.mmdd already exists
```

File setups have probably already run. Check the status of files, then run setups manually.

**ERROR: /usr/adm/acct/nite/wtmp.mmdd already exists, run setup manually**

You must perform the SETUP step manually, because the daily wtmp file already exists.

**ERROR: wtmpfix errors see /usr/adm/acct/nite/wtmperror**

wtmpfix detected a corrupted wtmp file. Refer to the section “Fixing Corrupted Files” in this chapter for details on fixing wtmp errors.

**ERROR: connect acctg failed: check /usr/adm/acct/nite/log**

acctcon1 encountered a bad wtmp file. Again, refer to the section “Fixing Corrupted Files” in this chapter for information on how to fix the file.

**ERROR: Invalid state, check /usr/adm/acct/nite/active**

the file statefile is probably corrupted. Check statefile and read active before restarting.

## Restarting runacct

runacct is normally run via cron only once per day. However, if an error occurs while executing runacct (as described above), it may be necessary to restart runacct. runacct has the following syntax:

```
runacct [mmdd [state]]
```

When called without arguments, runacct assumes that it is being invoked for the first time on the current day; this is how runacct is invoked by cron. The argument *mmdd* is necessary if runacct is being restarted and specifies the month and day for which runacct will rerun the accounting. The entry point for processing is based on the contents of statefile. To override statefile, include the desired entry *state* on the command line.

For example, to start `runacct`, you would enter:

```
$ nohup runacct 2> /usr/adm/acct/nite/fd2log &
```

To restart `runacct` on the 26th day of March:

```
$ nohup runacct 0326 2> /usr/adm/acct/nite/fd2log &
```

To restart `runacct` at state `WTMPFIX` on June 1st:

```
$ nohup runacct 0601 WTMPFIX 2>/usr/adm/acct/nite/fd2log &
```

All the above examples were run in the background (`&`) and use the `nohup` command so the process continues running even though you may log out.

## Daily Reports

`runacct` generates five basic reports upon each invocation. Brief descriptions of each report follow. Detailed descriptions of the reports are found in the following section, “Displaying `runacct` Reports—`prdaily`.”

- **Daily Line Usage Report**—summarizes connect session accounting since the last invocation of `runacct`. It provides a log of system reboots, power failure recoveries, and any other records dumped into `/etc/wtmp` via `acctwtmp`. In addition, it provides a breakdown of line utilization.
- **Daily Resource Usage Report**—gives a summary of resource usage per individual user: it basically merges all the total accounting records for individual users and displays the records, one per user.
- **Daily Command Summary**—summarizes resource usage data for individual commands since the last invocation of `runacct`. The data included in this report is useful in determining the most heavily used commands; you can use these commands’ characteristics of resource utilization when “tuning” your system.

This report is sorted by `TOTAL KCOREMIN`, an arbitrary but useful yardstick for calculating “drain” on a system.

- **Monthly Total Command Summary**—This report is exactly the same as the Daily Command Summary, except that the Daily Command Summary contains command summary information accumulated only since the last invocation of `runacct`, while the Monthly Total Command Summary summarizes commands from the start of the fiscal period to the current

date. In other words, the monthly report reflects the data accumulated since the last invocation of `monacct`.

- **Last Login**—simply gives the date each user last logged in to the system. This could be a good source for finding likely candidates for the archives, or getting rid of unused login directories.

## Displaying runacct Reports – `prdaily`

As `runacct` finishes executing, it deposits a report of the current day's accounting in the file `/usr/adm/acct/sum/rptmdd`, where `mdd` is the month and day that the report was generated. The `prdaily` command is used to display the contents of any daily report file created by `runacct`. Its syntax is:

```
prdaily [-l] [-c] [ mdd ]
```

where:

- `mdd` is an optional report date. If no date is specified, `prdaily` produces a report of the current day's accounting information. Previous days' accounting reports can be displayed by using the `mdd` option and specifying the exact report date desired.
- The `-l` option prints a report of exceptional usage by login name for the specified date. This option is used to determine which users are consuming excessive amounts of system resources. The limits for exceptional usage are kept in the file `/usr/lib/acct/ptelus.awk` and can be edited to reflect your installation's requirements.
- Valid only for the current day's accounting, the `-c` option is used to get a report of exceptional resource usage by command. This option is used to determine which commands are using excessive amounts of system resources. The limits for exceptional usage are maintained in the file `/usr/lib/acct/ptecms.awk` and can be edited to reflect your system's needs.

The reports produced by `runacct` were described briefly in the previous subsection. Now the reports are discussed in more detail.

## Daily Line Usage Report

In the first part of this report, the FROM/TO banner should alert you to which period is being reported. The times are the date-time that the last report was generated by `runacct`, and the date-time that the current report was generated. It is followed by a log of system reboots, shutdowns, power failure recoveries, and any other records dumped into `wtmp` by the `acctwtmp` command.

The second part of the report is a breakdown of line utilization. The **TOTAL DURATION** shows how long the system was in a multi-user state. The columns of the report are defined as follows:

Column	Description
LINE	The terminal line or access port being reported on.
MINUTES	The total number of minutes that the line was in use during the accounting period.
PERCENT	The percentage of <b>TOTAL DURATION</b> that the line was in use: $\text{PERCENT} = (\text{MINUTES} / \text{TOTAL DURATION}) * 100$
# SESS	The number of times that this port was accessed for a login session.
# ON	Historically, this column displayed the number of times that the port was used to log a user on; but since <code>login</code> can no longer be executed explicitly to log in a new user, this column should be identical to # SESS.
# OFF	This column reflects not only the number of times a user logged off, but also any interrupts that occurred on the line. Interrupts occur on a port when <code>getty</code> is first invoked. A <code>getty</code> is invoked when the system is brought to run-level 2. This column comes into play when # OFF exceeds # ON by a large factor. This usually indicates that the multiplexer, modem, or cable is going bad, or that there is a bad connection somewhere. The most common cause of this is an unconnected cable dangling from the multiplexer.

During real time, `wtmp` should be monitored as this is the file from which connect session accounting is taken. If it grows rapidly, execute `acctcon1` to determine which line is the noisiest. If the interrupting is occurring at a high rate, general system performance will be affected.

## **Daily Resource Usage Report**

This report gives a by-user breakdown of system resource usage. The format of this report is the same as that produced by the `prtacct` command. (For definitions of the data and format of this report, refer to the discussion of `prtacct` in the “Displaying Total Account Records – `prtacct`” section of this chapter.)

## **Daily and Monthly Command Summary**

These two reports are the same, except that the Daily Command Summary reports information only for commands executed since the last invocation of `runacct`; the Monthly Command Summary contains information on commands executed since the last invocation of `monacct`.

The output of this report is identical to that produced by `acctcms`. For definitions of the data found in this report, refer to the discussion of `acctcms` in the “Process Accounting” section of this chapter.

## **Last Login**

This report simply shows the last date and time that each user logged in. The longer it has been since a particular user logged in, the more likely it is that the user’s files could be archived, or maybe even that the user could be removed from the system.



## Creating Monthly Accounting Reports – monacct

`monacct` creates monthly summary files and reports; the resulting output is stored in the directory `/usr/adm/acct/fiscal`. After creating its monthly reports, it removes the old daily accounting files from the directory `/usr/adm/acct/sum` and replaces them with new summary accounting files.

`monacct` should be invoked once each month or accounting period. Its syntax is

```
monacct number
```

where *number* indicates which month or period it is (01=January, 12=December). If *number* is not specified, `monacct` assumes that it is being invoked for the current month; this default is useful if `monacct` is executed via `cron` on the first day of each month (as described in the “Daily Usage and Installation” section of this chapter).

Descriptions of the files created in the `acct/fiscal` directory follow:

- `cms?`—contains the total command summary file for the accounting period denoted by `?`. The file is stored in internal summary format. Therefore, to display this file, you must use the `acctcms` command. The following example shows how to display this file for the month of June:

```
$ acctcms -a -s /usr/adm/acct/nite/fiscal/cms06
```

- `fiscrpt?`—contains a report similar to that produced by `prdaily`. The report shows line and resource usage for the month represented by `?`. The following would display the fiscal accounting file for the month of November:

```
$ cat /usr/adm/acct/nite/fiscal/fiscrpt11
```

- `tacct?`—is the total accounting file for the month represented by `?`. To display this file, you must use the `prtacct` command. The following would display the total accounting summary file for the month of January:

```
$ prtacct /usr/adm/acct/fiscal/tacct01 "JANUARY TOTAL ACCOUNTING"
```

---

## Updating the Holidays File

The file `/usr/lib/acct/holidays` contains the information that System Accounting needs to distinguish between prime and non-prime time. It contains the following information:

- **Comment Lines**—Comment lines are entered by placing an asterisk (\*) as the first character in the line; they may appear anywhere in the file.
- **Year Designation Line**—This line should be the first non-comment line in the file and must appear only once. The line consists of three four-digit numbers (leading blanks and tabs are ignored). The first number designates the year; the second denotes the time (in 24-hour format) that prime time starts; the third gives the time that prime time ends and non-prime time starts.

For example, to specify the year as 1985, prime time at 9:00 a.m., and non-prime time at 4:30 p.m., the following entry would be appropriate:

```
1985  0900  1630
```

A special condition allowed for in the time field is that 2400 is automatically converted to 0000.

- **Company Holiday Lines**—These entries follow the year designation line. Company holidays are days when few people should be using the computer. Therefore, System Accounting assumes that non-prime time is in effect during the entire 24 hours of a specified holiday.

Company holiday lines have the following format:

```
day_of_year  Month  Day  Description of Holiday
```

The *day\_of\_year* field is a number in the range 1 through 366, corresponding to the day of the year for the particular holiday (leading blanks and tabs are ignored). The remaining fields are simply commentary and are not used by other programs.

---

### Note

As delivered, the `holidays` file contains valid entries for Hewlett-Packard's prime/non-prime time, and holidays. You should check this file and edit it as necessary to reflect your organization's requirements.

---

---

## Fixing Corrupted Files

System Accounting files may become corrupted or lost. Some of these files can simply be ignored or restored from the files saved through backup procedures. However, certain files must be fixed in order to maintain the integrity of System Accounting. Two of the files that must be fixed are `/etc/wtmp` and `/usr/adm/acct/sum/tacct`.

### Fixing wtmp Errors

The `wtmp` files seem to cause the most problems in the daily operation of System Accounting. When the date is changed and HP-UX is switched into multi-user mode, a set of date change records is written into `/etc/wtmp`. The `wtmpfix` command is designed to adjust the time stamps in the `wtmp` records when a date change is encountered. However, some combinations of date changes and reboots won't be caught by `wtmpfix` and cause `acctcon1` to fail. The following steps show how to "patch" a damaged `wtmp` file.

```
$ cd /usr/adm/acct/nite
$ fwtmp <wtmp.mmdd>wtmp.temp
    Using an editor, delete corrupted records or
    delete all records from beginning up to the date change
$ fwtmp -ic <wtmp.temp >wtmp.mmdd
$ rm wtmp.temp
```

If the `wtmp` file is beyond repair, create a null `wtmp` file. This will prevent any charging of connect time. `acctprc1` will not be able to determine which login owned a particular process, but it will be charged to the login that is first in the password file for that user ID.

## Fixing tacct Errors

If your installation is using System Accounting to charge users for system resource usage, the integrity of `sum/tacct` is quite important. If `sum/tacct` ever becomes corrupted, then check the contents of `sum/tacctprev` with the command `prtacct`. If it looks correct, then the latest `sum/tacct.mmdd` should be patched up, and `sum/tacct` should then be recreated. A simple patch procedure would be:

```
$ cd /usr/adm/acct/sum
$ acctmerg -a -v <tacct.mmdd>tacct.temp
    Using an editor, remove the bad records and
    write duplicate UID records to another file
$ acctmerg -i <tacct.temp >tacct.mmdd
$ acctmerg tacctprev <tacct.mmdd>tacct
$ rm tacct.temp
```

Remember that `monacct` removes all the `tacct.mmdd` files; therefore, `sum/tacct` can be recreated by merging these files together.

---

## Sample Accounting Shell Scripts

### **grpdu<sup>sg</sup>**

This shell script displays disk space usage totals for the users who are members of a specified group. The syntax of this command is:

```
grpdusg group_name
```

where *group\_name* is the name of the group for which disk space accounting information is to be generated.

For example,

```
$ grpdusg pseudo
```

generates disk space usage information for all the users in the group **pseudo**.

## The Shell Script

```
# Check for the group-name parameter.
#
if      [ $# -ne 1 ]
then    echo "\nUsage:  grpdusg group-name\n"
        exit 1
fi
echo    "\nOne moment please ... \n"
#
# Use the find command to find all the files whose owners are members of
# group-name.  Pipe the output from find into acctdusg which will accumulate
# disk space usage information for the users in group-name.
# NOTE:
#     - accounting data is temporarily stored in _${1}_tmp
#     - error messages are stored temporarily in _${1}_err
#     - if files exist that have no owners, then the names of
#       these files are stored in _no_owners
#
fn=_${1}_
find / -group $1 - hidden -print 2>${fn}err |acctdusg -u _no_owners >${fn}tmp
#
# Remove the _no_owners file if its size is not greater than zero.
#
if      [ -s _no_owners ]
then    echo "\nFiles having no owners exist--check _no_owners\n"
else    rm _no_owners
        echo "\nAll files have owners-- _no_owners not created\n"
fi
#
# Use echo and awk to display disk usage totals for this group.
#
echo "\nDisk space usage information (group is ${1}): \n"
awk 'BEGIN {print "\n_UID___USER
NAME_____BLOCKS"}
     { sum += $3 ;           # add up total disk blocks used
       print $0             # display information for user
     }
     END { print "\nTOTAL DISC SPACE USAGE= ", sum, "blocks" }' ${fn}tmp
#
# Remove temporary files, then exit.
#
rm ${fn}*
```

## **acct\_bill**

`acct_bill` takes as input a total accounting file and produces as output billing totals for all users found in the input file. The syntax of `acct_bill` is:

```
acct_bill [ mmdd ]
```

If the optional *mmdd* is not specified, then `acct_bill` takes as input the current day's total accounting file (`acct/nite/daytacct`); if *mmdd* is given, then input is taken from the total accounting file for the date specified by *mmdd* (`acct/sum/tacctmmdd`). Output is written to the file `billsmmdd`, where *mmdd* is the date given with the command, or the current date if *mmdd* was not specified with the command.

## **Examples**

To generate billing information for the current day, simply enter:

```
$ acct_bill
```

and the billing information will be stored in the file `acct/sum/billsmmdd`, where *mmdd* is the current date.

To create billing information for January 23rd, you would enter:

```
$ acct_bill 0123
```

after which the billing information would be stored in the file called:

```
acct/sum/bills0123
```

To automatically generate daily billing totals for all users, you should call `acct_bill` without the date argument from the `USEREXIT` state of `runacct`.

## **Output Produced by acct\_bill**

The output of `acct_bill` contains one line per user and has the following format:

```
user_ID      user_name      billing_amount
```

where *user\_ID* and *user\_name* identify the user who is being billed, and *billing\_amount* shows the total amount that the user is to be charged.

*billing\_amount* is computed by multiplying *accounting coefficients* (found in the shell script) by columns of the report generated by `prtacct`. Assuming that billing amounts are in dollars, the coefficients (as they are shown in the shell script that follows) produce the following billing amounts:

- ten cents for every minute of prime CPU time consumed
- five cents for every minute of non-prime CPU time consumed
- a half cent for every prime kcore minute used
- two-tenths of a cent for every non-prime kcore minute
- a half cent for every prime connect time minute
- two-tenths of a cent for every non-prime connect minute
- two-and-a-half cents for every block of disk space used
- two-and-a-half cents for every process spawned by the user
- ten cents for every connect session
- each fee unit charged via `chargefee` counts as one cent

You should experiment with this command by altering the coefficients to see how *billing\_amount* is affected. After gaining confidence with this shell script, you can alter the coefficients to suit your installation's needs.



## The Shell Script

```
_date='date +%m%d'
_outfile=/usr/adm/acct/sum/bills
_infile=/usr/adm/acct
#
# Set _infile and _outfile, based on whether or not mmdd was given
#
if      [ $# -eq 0 ]
then    # Generate billing data for current day.
        _infile=${_infile}/nite/daytacct
        _outfile=${_outfile}${_date}
else    # Create billing data for date given (mmdd).
        _infile=${_infile}/sum/tacct${1}
        _outfile=${_outfile}${1}
fi
#
# Create a file containing the ASCII equivalent of the input total
# accounting file (tacct_ASC.tmp_). The file can then be supplied as input
# to awk, which will generate billing data for each user.
#
acctmrg -a -t <${_infile} >tacct_ASC.tmp # output TOTAL amount first
acctmrg -a <${_infile} >>tacct_ASC.tmp # append users' total accounting records
#
# Using awk, compute billing totals for each user in the total accounting file.
#
awk 'BEGIN {
    # *****
    #   A C C O U N T I N G   C O E F F I C I E N T S
    # *****
    cpu_P =0.10 # 0.10 monetary units per minute of prime CPU time
    cpu_NP=0.05 # 0.05 monetary units per non-prime CPU minute used
    kcm_P =0.005 # for prime kcore minutes consumed
    kcm_NP=0.002 # for non-prime kcore minutes used
    con_P =0.005 # prime connect (real) time
    con_NP=0.002 # non-prime connect time used
    blk = 0.025 # number of blocks used
    prc = 0.025 # number of processes spawned
    ses = 0.10 # number of connect sessions
    fee = 0.01 # 100 charge units per monetary unit
    # *****
}
# Start computing billing amounts for each user.
{ _sum = cpu_P*$3 + kcm_P*$5 + con_P*$7          # compute prime usage
  _sum+= cpu_NP*$4+ kcm_NP*$6+ con_NP*$8        # add non-prime usage
```

```
_sum+= blk*$9 + prc*$10 + ses*$11 + fee*$13 # add remaining amounts
printf "%-8s %-10s %10.3f\n", $1, $2, _sum # display user total
}' tacct_ASC.tmp_ >$_outfile # write output from awk to appropriate file
rm tacct_ASC.tmp_ # remove the temporary ASCII file
```

---

## System Accounting Files

This section contains descriptions of the different files processed by HP-UX System Accounting. The files are grouped according to the directory in which they are found. In an HP-UX cluster, the `/usr/adm` directory is a CDF.

### Files in the `/usr/adm` Directory

Filename	Contents
<code>diskdiag</code>	Diagnostic output from the execution of disk space accounting commands.
<code>dtmp</code>	Output from the <code>acctdusg</code> program.
<code>fee</code>	Output from the <code>fchargefee</code> command (ASCII total accounting records).
<code>pacct</code>	The current active process accounting file.
<code>pacct?</code>	Process accounting files switched via <code>turnacct</code> switch.

## Files in the `/usr/adm/acct/nite` Directory

Filename	Contents
<code>active</code>	Used by <code>runacct</code> to record progress. It contains warning and error messages. <code>activemmdd</code> is the same as <code>active</code> after <code>runacct</code> detects an error.
<code>ctacct.mmdd</code>	Total accounting records created from connect session accounting where <code>mmdd</code> is the month and day the file was created.
<code>ctmp</code>	Output of <code>acctcon1</code> —connect session records.
<code>daycms</code>	ASCII daily command summary used by <code>prdaily</code> .
<code>daytacct</code>	Total accounting records for current day.
<code>disktacct</code>	Total accounting records created by the <code>dodisk</code> command.
<code>fd2log</code>	Diagnostic output from the execution of <code>runacct</code> (refer to <code>crontab</code> entry).
<code>lastdate</code>	The last day that <code>runacct</code> was executed, in <code>date #+%m%d</code> format. (Refer to <code>date(1)</code> for a description of <code>+%m%d</code> date format.)
<code>lock</code> and <code>lock1</code>	Used to control serial use of <code>runacct</code> .
<code>lineuse</code>	Terminal ( <code>tty</code> ) line usage report used by <code>prdaily</code> .
<code>log</code>	Diagnostics output from <code>acctcon1</code> .
<code>logmmdd</code>	Same as <code>log</code> after <code>runacct</code> detects an error.
<code>reboots</code>	Contains beginning and ending dates from <code>wtmp</code> , and a listing of reboots.
<code>statefile</code>	Used to record the current state being executed by <code>runacct</code> .
<code>tmpwtmp</code>	<code>wtmp</code> file, corrected by <code>wtmpfix</code> .
<code>wtmperror</code>	Error messages, if any, from <code>wtmpfix</code> .
<code>wtmperrormmdd</code>	Same as <code>wtmperror</code> after <code>runacct</code> detects an error.
<code>wtmp.mmdd</code>	The previous day's <code>wtmp</code> file.

## Files in the /usr/adm/acct/sum Directory

Filename	Contents
cms	Total command summary file for current month in internal summary format.
cmsprev	Command summary file without latest update.
daycms	Command summary file for previous day in internal summary format.
loginlog	Shows the last login date for each user.
pacct. <i>mmdd</i>	Concatenated version of all process accounting files for the date <i>mmdd</i> . This file is removed after reboot.
rpt <i>mmdd</i>	Daily accounting report for date <i>mmdd</i> .
tacct	Cumulative total accounting file for current month.
tacctprev	Same as <code>tacct</code> without latest update.
tacct <i>mmdd</i>	Total accounting file for date <i>mmdd</i> .
wtmp. <i>mmdd</i>	Saved copy of wtmp file for date <i>mmdd</i> . Removed after reboot.

## Files in the /usr/adm/acct/fiscal Directory

Filename	Contents
cms?	Total command summary for month ? in internal summary format.
fiscrpt?	Report similar to <code>prdaily</code> for the month ?.
tacct?	Total accounting file for the month ?.



# Contents

---

<b>8. Trusted Systems</b>	
Components of System Security . . . . .	8-3
Physical Security . . . . .	8-3
Procedural Security . . . . .	8-4
System Security . . . . .	8-4
Identification . . . . .	8-4
Authentication . . . . .	8-4
Controlling Access to System Resources . . . . .	8-5
Audit System Usage . . . . .	8-5
Security Policy . . . . .	8-5
Converting to a Trusted System . . . . .	8-7
Trusted Computing Base (TCB) . . . . .	8-9
Introduction to the TCB . . . . .	8-9
Parts of the TCB . . . . .	8-9
Access Control Lists (ACLs) . . . . .	8-10
Auditing . . . . .	8-11
Auditing Concepts . . . . .	8-12
Maintaining the Audit System . . . . .	8-12
Operating the Audit System . . . . .	8-12
Enabling the Audit System . . . . .	8-13
Unique Audit Identification for Each User . . . . .	8-15
Selecting Audit Events . . . . .	8-16
Analyzing the Audit Information . . . . .	8-16
Self-Auditing for Trusted Process . . . . .	8-16
Procedure for Audit File Overflow . . . . .	8-17
Auditing Tasks . . . . .	8-20
Starting or Halting the Auditing System . . . . .	8-21
Selectively Auditing Users . . . . .	8-23
Auditing Selectively by Events . . . . .	8-24

Display the Auditing Information . . . . .	8-26
Audit Overflow Prevention and Handling . . . . .	8-28
Auditing in Subsystems and Applications . . . . .	8-30
Auditing in a HP-UX Cluster Environment . . . . .	8-30
Audit files . . . . .	8-30
Audit Control System Calls . . . . .	8-30
Audit and RFA . . . . .	8-30
Audit and NFS . . . . .	8-31
Auditing and Real Time HP-UX Systems . . . . .	8-31
The Password File . . . . .	8-32
Security Procedure Awareness . . . . .	8-33
Additional Changes to Your System to Improve Security	8-33
System Administration Security Practices . . . . .	8-34
Users . . . . .	8-34
File Protection . . . . .	8-35
Software . . . . .	8-35
Performing General Administration Tasks in a Trusted System . . . . .	8-37



## Trusted Systems

---

The S300 6.5 release of HP-UX includes various new features which can be used to increase system security. These features are intended to fulfill the requirements of Class C2 of the Trusted Computer System Evaluation Criteria (TCSEC) published by the National Computer Security Center (NCSC). To fulfill these requirements, you must properly convert the system to conform to a trusted system configuration, as described in the following sections. Although the 6.5 release isn't being formally evaluated, formal evaluation of these features by the NCSC is expected in the 7.0 release.

A **trusted system** is a computer, including its hardware and software, that has a specific security rating according to guidelines set up by the U.S. government. HP-UX trusted system features include (all documented in this chapter):

- A means for converting to a trusted system configuration.
- A trusted computing base (TCB).
- Access Control Lists (ACLs).
- Auditing.
- An additional “secure” password file.
- Awareness of security procedures—both on the user level (as documented in the *Beginner's Guide to HP-UX*) and on the system administrator level (as documented in this manual).

This chapter describes the above security features, describes some general security concepts and terms, shows how to convert to a trusted system, and gives some additional security guidelines.

Specific security guidelines for system administration tasks are integrated throughout this manual, and are highlighted in notes such as:

---

**Security**      This is an example security note.

---

In the 6.5 revision of HP-UX, a trusted system configuration does not include any networking. This means we do not plan to submit our networking products for evaluation. However, you can still use our security features with networking products. For example, auditing, ACLs, and the secure password file work in a clustered environment, but the clustered environment will not be evaluated.

---

## Components of System Security

The components of computer system security are:

- Physical security—control of access to system components.
- Procedural security—establishment and control of security procedures.
- System security—control of access by the security features provided by the HP-UX operating system.
- Security policy.

### Physical Security

Physical security involves the prevention of physical damage to system hardware and the corruption of software. The term **hardware** includes the central processor unit (CPU), system console, terminals, and other peripherals, such as printers, disk drives, and tape drives. The term **software** includes the operating system, programs, and data.

The causes of damage to hardware and software range from deliberate sabotage or vandalism to inadvertent damage caused by unskilled users. Regardless of the cause, such damage usually can be prevented by restricting physical access to hardware and logon access to software.

Physical access to hardware is usually implemented with perimeter controls, which restrict entry into areas in which computer equipment is located. Perimeter controls include locked computer rooms, fenced building sites, and guard stations at building entrances. Physical access can be controlled by issuing keys and ID badges only to authorized persons.

Access to software is usually controlled by logon restrictions. Such restrictions include the use of passwords, establishment of accounts and groups, and control of user capabilities. Access to programs and files is provided by assigning users to accounts and enforcing the use of passwords. The physical aspect of securing access to software involves prevention of physical access to terminals, and limitations on or prevention of access via communication lines.

## **Procedural Security**

Procedural security deals with the establishment and enforcement of security procedures. Some of these procedures may be independent of the type or types of computers involved. Others may not. For example, perimeter security controls are usually similar for all types of systems. But desktop computers may require forms of antitheft protection not required by mainframes.

Procedural security regulates the performance of duties associated with system operation and use, and with the physical storage of system information. Common security practices include partitioning computer operating duties, rotating operators, and storing backup tapes at bonded, offsite depositories. Procedural security also encompasses and may regulate company policies that deal with information security, such as policies that regulate the way individuals manage their own passwords.

## **System Security**

System security is provided by security features built into HP-UX, by the ways in which the account structure of the system is organized, and by the auditing system. System security features fall into five categories:

- Identification of users.
- Authentication of users.
- Control of access to system resources.
- Auditing system usage.

### **Identification**

Every user must have a unique login identity, or ID, by which the user is identified as a legitimate system user. Without a valid ID, a user cannot log on to the system.

### **Authentication**

When a user logs on, the system attempts to authenticate the logon ID. The system checks the password file for the existence of the ID, then verifies the user's identity by checking the password (if required). Entry of an incorrect name or password or both is enough to prevent access to the system.

Although passwords are an optional feature of most systems, they are a must for any system that aspires to even the lowest levels of security.

## **Controlling Access to System Resources**

System performance can suffer if the system is overloaded with user processes. Setting limits on the number of processes that one user can run concurrently protects the system from inadvertent or deliberate attempts to degrade its performance.

## **Audit System Usage**

When activated, the auditing system maintains log records of system use. For example, the records can tell you how often abortive logons are attempted, and even identify the devices from which the attempts were made.

Refer to the section “Auditing” for more information.

## **Security Policy**

A computer security policy is a statement of rules that govern the behavior of system users, relative to system security.

A security policy covers the following aspects of computer operations:

- Types of facilities in which systems can be located.
- Who is allowed physical access to the system.
- Who is allowed to log on to the system.
- What capabilities are required to:
  - Use certain commands.
  - Use certain programs.
  - Access certain files.
  - Access certain devices.
- Which security features will be implemented (the following contains several examples of the types of facilities to be implemented):
  - Using ACLs to protect files and devices.
  - Requiring passwords.
  - Choosing and protecting passwords.

This list is not a comprehensive statement of security policy, but is a guide to what should be included in your security policy.

Once a security policy is developed, orientation and education becomes one of its most critical elements. Current and new users must be made familiar with the policy, and indoctrinated in its use. Periodic reinforcement of the policy message is a must to assure continuing compliance with the policy (and its updates).

---

## Converting to a Trusted System

To properly convert your system to a trusted system configuration, you must complete the following steps. These steps initialize the auditing subsystem and create the secure password file. You do not need to explicitly install or initialize ACLs because they are part of the basic HP-UX system:

1. You must have the AAUDIT fileset on your system.

If you loaded all filesets when you installed or updated to 6.5, then the AAUDIT fileset was loaded. To check, type:

```
ls /etc/filesets | grep AAUDIT
```

If the AAUDIT fileset is on your system, the system will respond by listing AAUDIT.

If the AAUDIT fileset is not on your system, the system will not list anything, and will give you a prompt. In this case, use the `update` program to load AAUDIT now.

2. Edit the `/etc/auditrc` script to activate and customize it.

At boot time the `/etc/rc` script calls the `/etc/auditrc` script, which initializes many things in the auditing subsystem.

The `/etc/auditrc` script is well-commented. To get your auditing system running, follow those comments, along with the additional notes here:

- You can set parameters in `audevent` to log different classes of events. Many different combinations of events and system calls can be audited. However, carefully consider your choices of events to log because if you log too many, two things can happen:
  - system performance can degrade—since you write an audit record for every audited event, it affects system performance.
  - file system space fills more rapidly—if you fill the audit file and an audit record can't be written, the operating system suspends all non-superuser processes not on the console. If a non-superuser process is on the console, it is logged off.

- The `audsys` and `audomon` commands use a concept called **switch points**. Switch points are thresholds where audit logging changes, or switches, from the primary to the alternate log file. The two switch points are:
    - Maximum audit file size.
    - Minimum remaining file system disk space.
3. Shut down your system using the `shutdown` command. For example, if you wish to shut down the system with no grace period for your users, type:

```
shutdown 0
```

4. Run the `tsconvert` (Trusted Systems Conversion) program by typing:

```
/etc/tsconvert
```

This program does the following:

- a. It creates a new secure password file that has just read-only permission for root.
  - b. It moves all encrypted passwords from the `/etc/passwd` file to the new secure password file, replaces the existing password field with “\*”, forces all users to use passwords, and creates a new audit ID for each user. This audit ID is a new field in the secure password file, and is used by the auditing subsystem. Refer to the section on “The Password File” later in this chapter for more information.
  - c. It enables auditing for all existing users by setting the auditing field to a 1.
5. Reboot the system by typing:

```
reboot
```

This reboot brings your system back to multiuser mode and starts the auditing system.



---

## Trusted Computing Base (TCB)

### Introduction to the TCB

The Trusted Computing Base (TCB) lies at the heart of a trusted system. The TCB contains all the elements of the computer system—hardware, firmware, and software—necessary to create a protective environment, able to provide user services required for a trusted computer system, and in combination, responsible for enforcing security policy and isolating objects (code and data) for which the protection is provided.

The TCB isolates resources to be protected so that they are subject to access control and auditing requirements. In HP-UX, the superuser is the TCB-privileged user. Security officers must be superusers.

In the HP-UX system, all login attempts are monitored. HP-UX gives as little information as possible to people attempting to access the system. You can keep track of login attempts through the auditing subsystem. At the C2 level, once a system itself is accessed legitimately, the user has access to information on the system for which the user has sufficient privilege.

### Parts of the TCB

The TCB consists of:

- The kernel and all key user programs, including login, getty, backup, restore, reconfig, and other utilities essential to run and support the system and maintain security.
- All security-related software, such as auditing.
- The Series 300 computer and all supported hardware.

---

## Access Control Lists (ACLs)

Access control lists are a key enforcement mechanism of discretionary access control (DAC) for specifying access to objects by users and groups more selectively than the traditional HP-UX DAC mechanism (permission bits) allows. DAC is discretionary because a subject can pass permissions (perhaps indirectly) to any other subject.

There is no system administration associated with ACLs because, as a form of DAC, they are manipulated by ordinary users. ACLs are an integral part of your HP-UX system. You cannot turn them on or off: they are available to all users at any time. However, if anyone on your system is using them, you must be aware of the following:

- ACL information is kept with files during backups only if you use the `fbbackup` command. All other archive commands issue a warning message, but do not back up ACL information.
- Do not put ACL information on the operating system files shipped on your install/update tape. HP will not support ACL information on these files when updating to future revisions. ACLs are appropriate only for user files and applications.
- ACLs are lost when copying files over a network connection such as RFA or NFS; however, when present, they control access over RFA or NFS.
- ACLs have full functionality in the HP-UX clustered environment; i.e., they behave the same as in the standalone environment.

For more information on access control lists, refer to the ACL manual entries, in particular *acl(5)* in the *HP-UX Reference Manual*. You can read about ACL commands for ordinary users in the *Beginner's Guide to HP-UX*.

---

## Auditing

Auditing fulfills part of the requirements for evaluated trusted systems. This section describes the functionality of the auditing subsystem.

In the HP-UX auditing system, you can audit many different types of information. For a complete list of events refer to the *audevent(1M)* and *audeventab* entries in the *HP-UX Reference*. Some examples of events you can audit are:

- Login—record physical location of user and all login attempts.
- File open—record name of file.
- Object deletion—record name of object (for example, files and IPC data objects).
- Operator/administrator actions.
- Other security relevant events.

The records in the audit file contain the following information:

- Date and time of the event.
- Identity of the user who directly caused the event.
- Identification of the event.
- Whether the outcome was a success or failure.
- If named object creation/deletion, then record the object name.
- If authorization and/or identification event type, then record the tty name.

The HP-UX audit system and its components are secured from tampering (as required for evaluation). The system meets the following criteria:

- The event recording mechanism is part of the TCB and is protected from unauthorized modification.
- The audit data itself is protected by the TCB from unauthorized access. The security officer (superuser) may access the audit data.
- The audit event on/off switches are part of the TCB and are tamper proof.

The HP-UX audit subsystem includes commands, system calls, and audit files. The files are protected by access control mechanisms. Only users with appropriate privileges (the superuser) may read, write, and execute the audit subsystem components.

## **Auditing Concepts**

This section contains conceptual information on the auditing system, and some guidelines on how to use the auditing system. For specific tasks within the auditing system, refer to the section “Auditing Tasks”.

This section contains information on:

- Maintaining the audit system.
- Operating the audit system.

## **Maintaining the Audit System**

As the system administrator or system security officer, it is your responsibility to maintain and protect the audit subsystem. You must initialize security audit procedures, including:

- Check the audit subsystem at least once a day.
- Review the audit log for unusual activities such as late hour logins, login failures, and failed accesses to system files.
- Quickly remove users who should no longer have access to the system.
- Archive the audit file as needed to prevent possible overflow.
- Periodically revise the currently selected events.
- Periodically revise the set of audited users.
- Do not follow any pattern or schedule for event selection and user selection.
- Constantly watch not only for external intrusion but also for possible internal violations of security policy.
- Retain the on-line audit file for 24 hours or less.
- Retain all audit trails stored off-line for a minimum of thirty days.
- Review the audit trail at least once a week.

## **Operating the Audit System**

All HP-UX installation tapes include the auditing system. You must decide whether or not to enable the audit system. The steps to enable the auditing system are covered in the section “Configuring a Trusted System”. This section

gives some conceptual information on enabling the audit system and its parts. There is information on:

- Enabling the Audit System.
- Unique Audit Identification for Each User.
- Selecting Audit Events.
- Analyzing the Audit Information.
- Self-Auditing for Trusted Processes.
- Procedure for Audit File Overflow.

**Enabling the Audit System.** The auditing system is enabled, and automatically restarted at bootup, when you follow the procedures in “Configuring a Trusted System”. By following those procedures you:

- Identify the current audit file.
- Identify the next audit file.
- Choose audit events.
- Configure the system to automatically start the auditing system and the `audomon` daemon, which monitors growth of the audit file.

Refer to the section “Procedure for Audit File Overflow” for a description of the current audit file, next audit file, free space switch point (FSS), and audit file switch point (AFS).

To support day to day administration of the auditing system, you can use the `audsys` command to:

- Turn the auditing mechanism on or off.
- Supply the current and next audit files for recording.
- Set the AFS values for these files.
- Manually switch audit files from the current to the next audit file.
- Disable the use of the next audit file.
- Display information about the status of the auditing system and the audit files; including the size of the audit files and the size of the filesystems on which they reside, as well as the percentage of the AFS size and filesystem space available.

`Audsys` uses the `audctl` system call to specify the current and next audit files to the kernel and maintains the names of these files and their AFS sizes in the `/.secure/etc/audnames` file. The `/.secure/etc/audnames` file contains two

records. The first record is the current audit file and AFS separated by a comma and the second is the next file name and AFS separated by a comma, as shown here:

```
current, AFS
next, AFS
```

`audsys` creates the file if it doesn't already exist, provided it is invoked with sufficient information to do so. Audit files specified to `audsys` for use by the auditing system must be nonexistent or regular empty files.

Note that in an HP-UX clustered environment, the audit files are automatically created as CDFs. Be aware of this when manipulating the files such as when copying or moving them.

If you use `audsys` to change the audit files while the auditing system is running, you must either:

- Specify the existing current and next files, as they exist in the `/.secure/etc/audnames` file.
- Specify the current file to be the next file that is in the `/.secure/etc/audnames` file.

To restart the auditing system using the names and AFS values already existing in `/.secure/etc/audnames`, use the `audsys` command with the `-n` option. This is how the auditing system is restarted, by default, using the `auditrc` script at bootup.

When using `audsys -n`, it chooses the current log file based on the following:

1. If the file specified in the second `/.secure/etc/audnames` record is empty, then `audsys` uses the first record for the current file.
2. If the file specified in the second `/.secure/etc/audnames` record is not empty, then `audsys` uses that file as the current file. That is because if the second record is a non-empty file, `audsys` assumes that a rollover (from current to next file) has taken place since the time the auditing files were last changed with `audsys`.

The `audomon` daemon uses the AFS values recorded in `/.secure/etc/audnames` to determine when the auditing system should switch from the current audit file to the next.

Refer to the section “Procedure for Audit File Overflow” for a detailed explanation of the audit file overflow procedure.

**Unique Audit Identification for Each User.** In a trusted system at the C2 level, users must be individually accountable for their security-relevant actions. This implies that a user should be accountable for all events in a session, from the login to the exit. Each existing user is assigned a unique audit identification number (audit ID) when you convert to a trusted system, and all new users are assigned an audit ID when you add them with `reconfig`. The audit ID is inherited by child processes (even when the user becomes another user via the `su` command), and can not be changed.

Use the `audusr` command to audit all users or selected individuals by user name. The `audusr` command with the add (`-a`) option sets the audit flag in the user field in the `/.secure/etc/passwd` file. When the specified user logs in, the system tags all processes associated with the user and begins auditing the actions of the user.

An audit record is generated when an audited user executes a currently enabled audited event, such as a `create` event. These selectable events are system-wide events. To stop auditing a user, run the `audusr` command with the delete (`-d`) option, which clears the audit flag in the password file. The user will not be audited at their next login.

You can also choose to audit all or none of the users who log in by running the `audusr` command with the `-A` option (for all the users) or `-D` option (no users).

When you add a new user to the system, the `reconfig` program automatically sets the audit flag in the password file.

Because process audit IDs, once set, do not change, do not invoke the following daemon processes from a login process:

```
cron
inetd
remshd (normally run from inetd)
rlogind (normally run from inetd)
rexecd (normally run from inetd)
rpc.rexd (normally run from inetd)
telnetd (normally run from inetd)
vt daemon
```

These daemons might attempt to set an audit ID, or call programs which do, and the attempts will fail. Do not use `inetd -k`.

**Selecting Audit Events.** This section describes some of the concepts in selecting audit events. For a description of the `audevent` command, and for examples, refer to the section “Auditing Selectively by Events”.

You can select which auditable events or auditable system calls should be logged in the audit file. Use the `audevent` command to choose the specific events or system calls that should be: not logged, logged if successful, logged if unsuccessful, or always logged no matter what the result was. Selectable events are event types such as `create`, `delete` and `open` (see the section called “Auditing Selectively by Events” or the `audevent(1M)` entry in the *HP-UX Reference*) and system calls. The auditable system calls are grouped into events (as documented in the `audevent(1M)`). When you select a particular event, you also select the system calls that correspond to that event.

Each call to `audevent` adds to previous calls. Note that if an event type affects both self-auditing commands and system calls and it is desired that they are set differently, the event type will have to be set first and then the individual system calls set or unset as appropriate. This is because by setting the event type, certain related system calls also get set automatically for the user.

Each auditable event has an event number that is used by auditing intrinsics and the `audisp` command. The event is stored in a highly compressed format and is not readable without the `audisp` command.

**Analyzing the Audit Information.** When you need to analyze the audit information, use the `audisp` command to analyze and display the audited information contained in the audit logfile. The `audisp` command prints audit records optionally by user, tty, event name, system call, success/failure, and time. If multiple audit files are specified, the command merges the files into a single audit trail in order of the time stamp in each record. Refer to the section “Display the Auditing Information” for details.

**Self-Auditing for Trusted Process.** In addition to the above auditing features, the system allows trusted processes to perform their own auditing. Trusted processes may wish to suspend/resume standard auditing to save disk space and provide a summary record of what the process does. An example is the system administration package (`reconfig`) which writes an audit record showing that a new user was added to the system. With the features of self-auditing you can



obtain a higher level view of the events than would be given by the lower level kernel calls. This higher level audit record is a summary of what the trusted process has done and does not include some temporary opening and closing of intermediate files.

Each trusted process, when self auditing, performs the following sequence:

1. Call `audswitch` to suspend low level auditing.
2. Execute existing code.
3. Create the audit record.
4. Call `audwrite` with audit record.
5. Call `audswitch` to resume low level auditing.

The trusted commands that are self-audited are:

<code>login</code>	The login utility.
<code>reconfig</code>	The system administration package.
<code>init</code>	Change run levels, users logging off.
<code>audsys</code>	Start or halt the auditing system.
<code>audusr</code>	Select users to be audited.
<code>audevent</code>	Select events to be audited.
<code>audisp</code>	Display the audit data.
<code>at</code>	Execute command at a later time.
<code>crontab</code>	Execute commands on a regular basis.
<code>cron</code>	Clock daemon.
<code>newgrp</code>	Change effective group.
<code>chsh</code>	Change login shell.
<code>chfn</code>	Change gecos field of password file.
<code>passwd</code>	Change password.
<code>pwck</code>	Password/group file checker.

**Procedure for Audit File Overflow.** When the audit system is enabled, a **current audit file** must be specified along with an optional backup file (called the **next audit file**). As audit records are written into the audit file, the sizes of the file and the file system that it is on need to be monitored closely to prevent an overflow of the file space to insure integrity of audit data. The mechanism for audit file overflow prevention/handling is described in this section.

When the current and next audit files are specified, a “switch point” size in kilobytes is required for each file. This switch point, called the **audit file switch** (AFS), is defined to be the size of the audit file at which point the audit system will attempt to switch to the next audit file (if available). If the next audit file is

not available, the attempt fails and a warning message is sent, but auditing still continues with the current audit file.

Separate from the AFS is a parameter called **free space switch** (FSS). Just as the AFS serves as a control parameter for the audit file, the FSS acts as a control for the file system on which the audit file resides (henceforth referred to as the audit file system). The FSS is defined to be the percentage of free space left on the audit file system at which point a switch to the next audit file is attempted. Again, the attempt may fail due to the absence of a next audit file, in which case a warning message is sent and auditing continues with the current audit file, unless the file system free space limit is reduced. In this case audit writes, as well as ordinary writes, fail.

Note that the AFS and FSS are two independent parameters, one for the audit file, and the other for the file system. Depending on the type of activities on the system and how the parameters are set, either of them can be reached before the other. Both the AFS and FSS cause the same action: to attempt to switch to the next audit file.

If you have more than one file system on your system, create the current and next audit files on different file systems. If they are on the same file system, and the file system fills (triggering the FSS), you will have the same space problem as before.

The sizes of both the audit file and the audit file system are monitored periodically by a daemon called `audomon` (which stands for AUDit Overflow MONitor), which is typically spawned as part of the `init` process when the system is booted. Its primary function is to check the audit file and the file system against the AFS and FSS and, if either is reached, to automatically switch to the next audit file (if it exists).

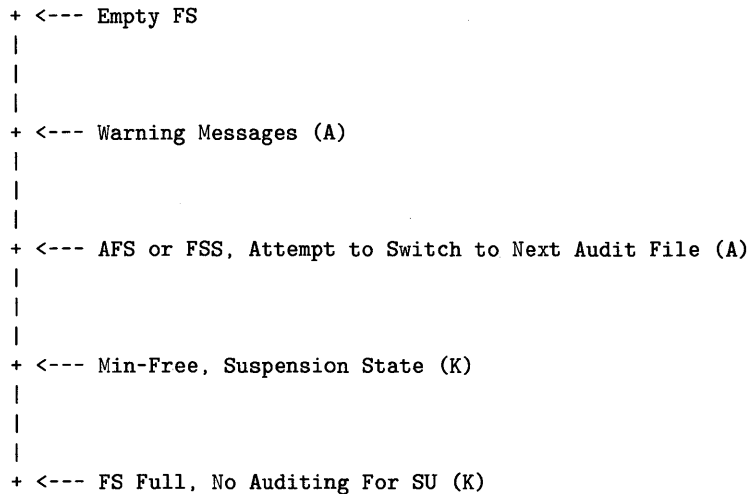
When invoking `audomon` you can specify how often warnings are sent when the audit system approaches the switching points. Refer to the section “Audit Overflow Prevention and Handling” for more details.

What happens if the attempt to switch over fails, and the audit file space continues to be used up? When the audit file system reaches `min-free` (in file system terms), the system enters into a “suspension state”. In the suspension state, all non-privileged processes are prohibited from writing to the applicable file system. Because of this, all non-privileged processes (and users) that are being audited are suspended if and when they require a write of auditing information.

Only the superuser, logged in at the console, is allowed to continue working. The system suspends all non-superuser processes, and logs off any non-superuser logged in at the console. Other users can log in, but are suspended immediately if they are not at the console, or they will be logged off immediately if they try to log in at the console. This situation lasts until new empty audit files are specified with the `audsys` command.

If, during suspension state, the space problem continues to worsen due to improper actions taken by the privileged user(s), the entire file system may eventually become full. Auditing then stops for the privileged users, who can continue to do any operations as normal, but no audit records are written. Note that the non-privileged users are still suspended.

The following graph depicts the various possible stages of the audit system and the corresponding actions taken.



Note:

- (A) refers to actions taken by the daemon
- (K) refers to actions taken by the kernel.

## **Auditing Tasks**

The following pages describe the tasks, and the commands with which you perform the task, in the auditing subsystem. These are:

- Starting or Halting the Auditing System
- Selectively Auditing Users
- Auditing Selectively by Events
- Displaying the Auditing Information
- Audit Overflow Prevention and Handling

All of the auditing commands are restricted to users with appropriate privilege.

---

## Starting or Halting the Auditing System

Start or halt the auditing system and set or display audit file information.

### Command Syntax

```
audsys [ -nf ] [ -c file -s cafs ] [ -x file -z xafs ]
```

Refer to the *audsys(1M)* entry in the *HP-UX Reference* for a description of the options.

### Description

Audsys lets you start or halt the auditing system; specify the auditing system “current” and “next” audit files (and their switch sizes); or display auditing system status information. This command is restricted to users with superuser privilege.

The “current” audit file is the file to which audit records are written by the auditing system. When the “current” file grows to its Audit File Switch (AFS) size or its File Space Switch (FSS) size the auditing system switches to write to the “next” audit file. The auditing system switches audit files by setting the “current” file to what was the “next” file and setting the new “next” file to null. The “current” and “next” files may reside on different file systems. In fact, the two files *should* reside on different file systems if possible.

When invoked without arguments, *audsys* displays the status of the auditing system. This status includes information describing whether auditing is on or off, the names of the “current” and “next” audit files, and a table listing their switch sizes and the sizes of filesystems they are on, as well as the space available in terms of a percentage of the switch sizes, and filesystem sizes.

### Files

*/.secure/etc/audnames*—File maintained by *audsys* containing the “current” and “next” audit file names and their switch sizes.

## Starting or Halting the Auditing System

### Examples

1. Start the auditing system with `/mnt/a1/auditfile1` with an AFS of 10000 Kbytes as the current audit file and `/mnt/a2/auditfile2` with an AFS of 20000Kbytes as the next audit file.

```
audsys -n -c /mnt/a1/auditfile1 -s 10000 -x /mnt/a2/auditfile2 -z 20000
```

2. Halt the auditing system.

```
audsys -f
```

3. Restart the auditing system, resuming use of the audit files previously in use.

```
audsys -n
```

4. Change the current audit file to `/mnt/a3/auditfile3` with an AFS of 30000 Kbytes.

```
audsys -c /mnt/a3/auditfile3 -s 30000
```

5. Setting the current file to the previous next file forces the auditing system to switch to the next audit file immediately. Invoking `audsys` as demonstrated, when the next file for this example is `/mnt/a2/auditfile2`, causes the auditing system to immediately begin using `/mnt/a2/auditfile2` as the current file.

```
audsys -c /mnt/a2/auditfile2 -s 10000
```

---

## Selectively Auditing Users

Select users to audit.

### Command Syntax

```
audusr [ [ -a user ] ... ] [ [ -d user ] ... ] [ -A | -D ]
```

Refer to the *audusr(1M)* entry in the *HP-UX Reference* for a description of the options.

### Description

`audusr` is used to specify *users* to be audited or excluded from auditing. Without arguments, `audusr` displays the command usage. This command is restricted to users with superuser privilege.

Users specified using `audusr` are audited (or excluded from auditing) starting at their next log-in session. Users logged in to the system when `audusr` is used are unaffected by the options specified; however, any users who log in after `audusr` is invoked are audited or excluded from auditing accordingly.

### Files

`/.secure/etc/passwd`—File containing flags to indicate if users are audited.

### Examples

1. Audit the users john, joan and wiz

```
audusr -a john -a joan -a wiz
```

2. Stop auditing the users john, jane and jim

```
audusr -d john -d jane -d jim
```

3. Audit everyone who logs into the system.

```
audusr -A
```

4. Do not audit any users.

```
audusr -D
```

---

## Auditing Selectively by Events

Change or display event or system call audit status.

### Command Syntax

```
audevent [ -P|-p ] [ -F|-f ] [ -E ] [ -e event ] ... [ -S ] [ -s syscall ] ...
```

Refer to the *audevent(1M)* entry in the *HP-UX Reference* for a description of the options and of the event types.

### Description

In the HP-UX auditing subsystem, events are grouped into event types to provide simpler control of auditing. Some event types map to groups of system calls with similar functions; others are for self-auditing events. Refer to the *audevent(1M)* entry in the *HP-UX Reference* for a list of selectable event types.

You can specify whether you want to log successes and/or failures for each event type. Because certain commands often make unnecessary closes on unopened files, the `close` event is generally not turned on for logging failures except in special cases.

`Audevent` changes the auditing status of the given events or system calls. Events specify either names associated with certain self-auditing commands, or names used to select related system calls.

Only the superuser may change or display the audit status. The results of this command take effect immediately.

### Examples

1. Audit the `modaccess`, `remove`, and `delete` events for both success and failures.

```
audevent -P -F -e modaccess -e remove -e delete
```

2. Remove `login`, `create`, and `open` events from the auditing list.

```
audevent -p -f -e login -e create -e open
```

3. Audit all selectable events. Log both success and failures.



## Auditing Selectively by Events

```
audevent -P -F -E
```

4. Do not audit any selectable events.

```
audevent -p -f -E
```

5. Audit all failures of the open and close system calls. (The names of the auditable system calls can be found in the *audevent(1M)* entry in the *HP-UX Reference*.)

```
audevent -F -s open -s close
```

6. Show the current events being audited.

```
audevent
```

Note that this is a system wide event audit such that any audited user triggering any of the enabled events will get audited. If all selectable events are disabled, only the mandatory events are in effect.

---

## Display the Auditing Information

Display the audit information as requested by the parameters.

### Command Syntax

```
audisp [-u username] [-e eventname] [-c syscall] [-p] [-f]
        [-l ttyid] [-t start_time] [-s stop_time] audit_filename ...
```

Refer to the *audisp(1M)* entry in the *HP-UX Reference* for a description of the options.

### Description

The *audisp* command analyzes and displays the audit information contained in the specified audit file(s). If multiple files are specified, it merges the files into a single audit trail in order of the time stamp of the audit records. The command allows the user to limit the information displayed by specifying the various options. This command is restricted to superusers.

---

**Note** If you are in an HP-UX clustered environment, you must use the following syntax:

```
audisp cdf_auditfile+/*
```

If you do not, you will see the information for one element only.

---

Any unspecified option is assumed to mean the most unrestricted specification. For example, a missing *-u username* option causes all users' audit information to be displayed, as long as it meets the specification by all other specified options. By the same principle, a *-t start\_time* option without a *-s stop\_time* displays all audit information starting from the *start\_time* to the end of the audit trail. The *audisp* command without any options displays all recorded information from the start of the audit file to the end.

## Examples

1. What events were audited from 3:10 pm, June 1 to 10:30 am, June 3?

```
audisp -t 06011510 -s 06031030 filename(s)
```

2. Did terminal tty05 have unsuccessful logins from Oct 11 through Oct 13?

```
audisp -l tty05 -e login -f -t 10110000 -s 10132400 filename(s)
```

3. What files were created by user jane from July 4, 1988 to July 10, 1988?

```
audisp -u jane -e create -t 0704000 -s 07102400 filename(s)
```

4. What files did user john attempt to delete unsuccessfully?

```
audisp -u john -e delete -f filename(s)
```

5. Did user john use tty05 on Dec. 25 1988?

```
audisp -u john -l tty05 -t 12250000 -s 12252400 filename(s)
```

---

## Audit Overflow Prevention and Handling

Audit overflow monitor daemon.

### Command Syntax

```
/usr/bin/audomon [ -p fs ] [ -t sp_freq ] [ -w warning ]
```

Refer to the *audomon(1M)* entry in the *HP-UX Reference* for a description of the options.

### Description

*Audomon* monitors the capacity of the current audit file and the file system where the current audit file resides, and prints out warning messages when either of them approaches full. It also checks the audit file and the file system against two switch points: **File Space Switch** (FSS) and **Audit File Switch** (AFS), and if either of them is reached, an automatic switch to the backup audit file takes place (if the backup file is available). The following is what the switch points mean.

**FSS**            The File Space Switch is specified as the amount of free space left as a percentage of the total disk space. When the free space of the file system reaches this percentage, *audomon* checks to see if a backup audit file is available and if so, it switches over and makes the backup the current audit file.

This percentage should be specified to be greater than the *min\_free* parameter of the file system to ensure that the switch takes place before *min\_free* is reached.

The default of FSS is 20(%).

**AFS**            The Audit File Switch is specified (with the *audsys* command) as the size of the audit file in Kbytes. When the size of the audit file (or in an HP-UX clustered environment, the total size of the CDF elements) reaches this point, *audomon* checks to see if a backup audit file is available and if so, it switches over and the backup becomes the current audit file.

## Audit Overflow Prevention and Handling

Note that this is a parameter of `audsys`, not of `audomon`. `Audomon` only checks it against the current size of the audit file and switches to the backup if necessary and possible. Please refer to the *audsys(1M)* entry in the *HP-UX Reference* for details on the specification of Audit File Switch.

If either of the switch points is reached but a backup file is not available, a warning message is printed.

`Audomon` is typically spawned as part of the `init` process when the system is booted.

### Files

`/.secure/etc/audnames`—contains the names of the current and next audit files, and their AFS switch points.

## Auditing in Subsystems and Applications

### Auditing in a HP-UX Cluster Environment

Although the HP-UX clustered environment will not be evaluated, Hewlett-Packard has insured all auditing commands work completely and correctly in an HP-UX cluster.

The HP-UX clustered environment presents a single file system hierarchy view. In a single system the administrator does not have to separately administer auditing on each system in the cluster. Auditing supports this single-system view.

**Audit files.** There would be a significant performance penalty associated with multiple cluster nodes writing to a single audit file. Hence, audit files are automatically created as context dependent files (CDFs). Each cnode has its own audit file. The NCSC has indicated that multiple synchronized audit trails are acceptable.

The `audisp` command provides the capability to merge multiple audit files into a single audit trail. To do this, you must use the following syntax:

```
audisp cdf_auditfile+/*
```

If you do not, you will see the information for one element only.

**Audit Control System Calls.** Audit functions that affect single files affects the entire cluster as a result of the single file system view. Thus, setting which users to audit works unchanged in the cluster environment.

The self-audit calls and setting the audit identifier affect only the current process. These are all done on the local processor.

Other functions affect global kernel variables. When the kernel state is modified a broadcast message updates the state on each node using the same mechanism as `mount` and `sync`. This happens when turning auditing on or off, specifying events to audit, specifying audit file names, and when the kernel rolls over to the next audit file.

### Audit and RFA

Although RFA is not part of our evaluated configuration, when a trusted HP-UX system is an RFA client, auditing works normally. The system calls that access files across RFA go through the normal system call handler on the client and

are audited appropriately. The server, however, does not audit accesses made by RFA.

## **Audit and NFS**

Although NFS is not part of our evaluated configuration, when a trusted HP-UX system is an NFS client, auditing works normally. The system calls that access files across NFS go through the normal system call handler on the client and be audited appropriately. The server, however, does not audit accesses made by NFS. Yellow pages and shadow passwords also don't work together fully. Refer to Appendix D of the *Using and Administering NFS Services* for more information.

## **Auditing and Real Time HP-UX Systems**

Auditing increases system overhead. Therefore, a system with real time needs may choose to select a minimum of selectable events. If you select fewer events you reduce the amount of system auditing. In other cases where the real time process does not execute audited events the system through-put may not be hindered.

You should balance the need for system security versus system performance. The kernel portion of the audit system meets HP-UX's real time requirements.

---

## The Password File

The Trusted System maintains two authentication files—`/etc/passwd` and `/.secure/etc/password`. Every system user has entries in both files, and login looks at both entries to authenticate login requests.

The `/etc/passwd` file contains descriptions of the accounts on your HP-UX system. The file is readable by all users, but only the superuser can modify it. As in non-trusted systems, each entry in the `/etc/passwd` file has seven fields, separated by colons. However, the encrypted password field contains a single asterisk (\*) instead of the encrypted password. A typical password file entry looks like this:

```
jyb:*:255:17:Joe Brown  :/users/jyb:/bin/sh
```

The `/.secure/etc/password` file contains the actual encrypted passwords, and is accessible (read-only) only to root. An example entry is:

```
jyb:a9x0EmtQsK6qc:214:1
```

The four fields of the `/.secure/etc/password` file contain the following information:

- The user name (or login name) consisting of up to eight characters.
- The encrypted password.
- Audit ID (used by the auditing subsystem at login time)
- Audit flag (1=on, 0=off)

You can use the routines in `getpwent`(LIBC) to gain access to the components of the secure password file as well as the components of the `/etc/passwd` file. Refer to the *HP-UX Reference* for details.



---

## Security Procedure Awareness

The above sections outline the required changes to convert to and use your HP-UX system in a trusted manner. This section lists some additional (optional) changes to help further improve security on your system, and discusses some general security practices for system administrators.

### Additional Changes to Your System to Improve Security

- Create the `/etc/securetty` file.

In this file you specify all `tty` device file names through which the superuser (`root`) can log in. This limits the number of terminals from which you can access superuser privileges (except by the `su` command). For more information on this file, refer to the *login(1)* entry in the *HP-UX Reference*.

- Modify the system login scripts to enforce security measures. Suggested modifications are:
  - Change the `umask` value in the system-wide shell startup files to allow read and write capability to owners only, for example `umask 077`.
  - Verify that the `PATH` environment variable includes only those directories where official HP-UX programs and commands reside. The `PATH` should not, for example, include the current directory as the first place to look. It also should not include null entries, which stand for the current directory.

Password aging is sometimes used to improve security. However, use discretion with a password aging policy as it can sometimes backfire. Because users are forced to change their password while logging in, they are in a hurry, they often do not choose good passwords, or simply switch between two passwords.

Some security measures you may have used in the past, for example the `btmp` file, are still available, but the new trusted system software now accomplishes the same function through auditing.

## System Administration Security Practices

The system is subject to a variety of potential threats, intentional or accidental, physical, or resulting from clever use of software.

### Users

Make sure users are on the system purposefully, on a need-to-know basis. Any unusual or irresponsible activity should be reported to authorities. This might include the presence of unaccounted-for programs or unexpected software behavior.

Since HP-UX is a relatively open system, access to software is ultimately controlled by login restrictions. Once a user gains access to the system, he or she has access to virtually all available functionality.

As the administrator you should preach user awareness to safeguard passwords. A login and password is the first and most important line of defense for access control. Users should choose passwords that are difficult to guess. Users can change their password with the `passwd` command.

Good passwords include:

- Misspelled words, such as `compoter` or `moon`.
- Mnemonic tricks, such as `t00l` or `billy`.
- Multi-word combinations, such as `tie_this_together` or `just-one`.
- Random strings (these are good to prevent breakins, but are hard to remember).
- Combinations of the above, such as `2-long`.

In general, don't choose anything obvious for a password. Bad passwords include:

- Proper names (even with number at end).
- Personal information—for example, Social Security Number, anniversary, birthday, or pet's name.
- Combinations of login name.
- English words.

Don't use a word found in a dictionary. Password-cracking programs can use dictionary lists. Don't even use a dictionary word with a number attached to it.

Tricks to help maintain good passwords:

- Change your password often.
- Don't store password in the function keys of a terminal.
- Do whatever is necessary to prevent `/etc/passwd` from being overwritten. `/etc/passwd` should be writable only by root.
- Use a password aging policy.
- Force the user to use passwords.

## File Protection

As shipped, the permissions for all core programs are set as securely as possible without being unnecessarily restrictive.

One way to enforce proper file and directory access permissions is to set the `umask` in the system-level login files (refer to the discussion on `umask` in *A Beginner's Guide to HP-UX*). The most restrictive `umask` is:

```
umask 077
```

which yields the following:

```
-rw-----          for files
-rwx-----          for directories and executables
```

## Software

Software restrictions include the use of passwords, establishment of accounts and groups, control of user capabilities, and use of access control lists. Access to security features should be restricted to those responsible for their use.

Threat by penetration can lead to destruction or compromise of programs and data. This could come about through covert channels, that is, information propagated deviously. You should use the auditing subsystem to monitor this.

There are many types of software penetration. Some are listed and defined here:

- Time bombs**     Once a system is penetrated, time bombs installed can “explode” later.
- Trap door**        A hidden software or hardware mechanism that permits system protection mechanisms to be circumvented. It is activated in some non-apparent manner (for example, by “random” key sequence at a terminal).

**Trojan horse** A computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the legitimate authorization of the invoking process to the detriment of security. For example, making a “blind copy” of a sensitive file for the creator of the Trojan Horse.

To help prevent software penetration by Trojan horses:

- Do not accept binary programs onto a trusted system.
- Don't put the current directory as the first directory in your PATH environment variable. If you do, you might inadvertently execute a non-TCB version of the program.
- Always type `/bin/su` to invoke the switch user command.
- Look for Trojan horse potential in archive programs—`shar`, `tar`, `cpio`.

**Viruses** Programs, attached to other programs, that propagate onto a system. They're often hidden in games and other public-access software. When root runs an infected program, the entire system can be infected.

You, as the system administrator or security officer, must watch for and diligently guard against all potential threats. Don't run or copy games or other public-access software or from a network into a secure environment. Know where you got your code.

---

## Performing General Administration Tasks in a Trusted System

While performing everyday system administration tasks, there is often a “trusted” method, or additional “trusted” reminders. Throughout the *System Administrator Manual*, there are special security notes to emphasize this “trusted” method of accomplishing a task. For your convenience the security note information is summarized here:

- In “Backing Up and Restoring the File System”:

If you are running a trusted system you must use the `fbackup` program for backups, and `frecover` for recovering files. This is the only backup and recovery method that preserves ACL information.

- In “Adding/Removing Users”:

If you are running a trusted system you must use the `reconfig` program to add new users. Do not add or remove users with the manual method.

- In “Shutting Down the System”:

Never leave your system in the system-administration run-level. This effectively runs as the root user so you have all privileges. Also, since no one has logged in, any auditing is meaningless.

- In “Configuring Device Drivers and I/O Cards Using `reconfig`”:

If you write your own device driver you invalidate the Trusted Computing Base. Use only HP-supplied device drivers in your kernel.

- In “Setting the System Clock”:

You must correctly set the date and time since the auditing subsystem depends on the system clock to log events.

- In “Creating and Using a Recovery System”:

If you are running a trusted system you must lock the recovery system tape. When you boot from a recovery system you are the superuser in an administrative mode. This means that you have all privileges. You also have no auditing when booted from the recovery system.

- In “Installing Optional Software and Updating Your HP-UX Kernel”:

To comply with trusted systems criteria you cannot *update* to 6.5, but must *install* 6.5.

# Contents

---

## A. Using the fsck Command

Introduction . . . . .	A-1
fsck Errors That May Occur . . . . .	A-5
Initialization Phase Errors . . . . .	A-6
“C” option? . . . . .	A-6
cannot alloc NNN bytes for “XXX” . . . . .	A-6
Can’t open checklist file: . . . . .	A-6
Can’t stat root . . . . .	A-6
Can’t stat ... Can’t make sense out of ... . . . . .	A-7
“file_system_name” is not a block or character device; OK . . . . .	A-7
Can’t open ... . . . . .	A-7
“file_system_name”: (NO WRITE) . . . . .	A-7
Other messages: . . . . .	A-8
INTERNAL INCONSISTENCY: “message” . . . . .	A-8
CAN NOT SEEK: BLK “bn” (CONTINUE) . . . . .	A-8
CAN NOT READ: BLK ... (CONTINUE) . . . . .	A-9
CAN NOT WRITE: BLK ... (CONTINUE) . . . . .	A-9
Phase 1 Errors: Check Blocks and Sizes . . . . .	A-9
CG ... : BAD MAGIC NUMBER . . . . .	A-10
UNKNOWN FILE TYPE I= ... (CLEAR) . . . . .	A-10
LINK COUNT TABLE OVERFLOW (CONTINUE) . . . . .	A-10
“block_number” BAD I= ... . . . . .	A-10
EXCESSIVE BAD BLKS I= ... (CONTINUE) . . . . .	A-11
“block_number” DUP I= ... . . . . .	A-11
EXCESSIVE DUP BLKS I= ... (CONTINUE) . . . . .	A-11
DUP TABLE OVERFLOW (CONTINUE) . . . . .	A-12
PARTIALLY ALLOCATED INODE I= ... (CLEAR) . . . . .	A-12
INCORRECT BLOCK COUNT I= ... (CORRECT) . . . . .	A-12

BAD INDIRECT ADDRESS: IND BLOCK n[m] = val	
I= ... (CORRECT) . . . . .	A-13
Phase 1b: Rescan for More Dups . . . . .	A-13
"block_number" DUP I= ... . . . . .	A-13
Phase 2: Check Path Names . . . . .	A-13
ROOT INODE UNALLOCATED. TERMINATING . . . . .	A-13
NAME TOO LONG . . . . .	A-14
ROOT INODE NOT DIRECTORY (FIX) . . . . .	A-14
DUPS/BAD IN ROOT INODE (CONTINUE) . . . . .	A-14
I OUT OF RANGE I= ... (REMOVE) . . . . .	A-14
UNALLOCATED I= ... (REMOVE) . . . . .	A-15
DUP/BAD I= ... (REMOVE) . . . . .	A-15
ZERO LENGTH DIRECTORY I= ... (REMOVE) . . . . .	A-15
DIRECTORY TOO SHORT I= ... (FIX) . . . . .	A-16
DIRECTORY CORRUPTED I= ... (FIX) . . . . .	A-16
BAD INODE NUMBER FOR '.' I= ... (FIX) . . . . .	A-16
MISSING '.' I= ... (FIX) . . . . .	A-17
MISSING '.' I= ... CANNOT FIX, FIRST ENTRY IN	
DIRECTORY CONTAINS ... . . . . .	A-17
MISSING '.' I= ... CANNOT FIX, INSUFFICIENT	
SPACE TO ADD '.' . . . . .	A-17
EXTRA '.' ENTRY I= ... (FIX) . . . . .	A-17
BAD INODE NUMBER FOR '..' I= ... (FIX) . . . . .	A-18
MISSING '..' I= ... (FIX) . . . . .	A-18
MISSING '..' I= ... CANNOT FIX, SECOND	
ENTRY IN DIRECTORY CONTAINS ... . . . .	A-18
MISSING '..' I= ... CANNOT FIX, INSUFFICIENT	
SPACE TO ADD '..' . . . . .	A-18
EXTRA '..' ENTRY I= ... (FIX) . . . . .	A-19
Phase 3: Check Connectivity . . . . .	A-19
UNREF DIR I= ... (RECONNECT) . . . . .	A-19
SORRY, NO lost+found DIRECTORY . . . . .	A-19
SORRY. NO SPACE IN lost+found DIRECTORY . . . . .	A-20
DIR I= ... CONNECTED. PARENT WAS I= ... . . . . .	A-20
BAD CONTINUATION INODE NUMBER. I= ...	
CLEAR? . . . . .	A-20
Phase 4: Check Reference Counts . . . . .	A-20
UNREF CONTINUATION INODE CLEAR? . . . . .	A-21



BAD LINK COUNT IN CONTINUATION INODE.	
I= ... (LINK COUNT SHOULD BE 1)	
CORRECT? . . . . .	A-21
UNREF FILE I= ... (RECONNECT) . . . . .	A-21
(CLEAR) . . . . .	A-21
SORRY. NO lost+found DIRECTORY . . . . .	A-22
SORRY. NO SPACE IN lost+found DIRECTORY . . . . .	A-22
LINK COUNT ... (ADJUST) . . . . .	A-22
UNREF ... (CLEAR) . . . . .	A-22
BAD/DUP ... (CLEAR) . . . . .	A-23
FREE INODE COUNT WRONG IN SUPERBLK	
(FIX) . . . . .	A-23
Phase 5: Check Cylinder Groups . . . . .	A-23
CG ... : BAD MAGIC NUMBER . . . . .	A-23
EXCESSIVE BAD BLKS IN BIT MAPS (CONTINUE)	A-24
SUMMARY INFORMATION "t" BAD . . . . .	A-24
"x" BLK(S) MISSING . . . . .	A-24
FREE BLK COUNT WRONG IN SUPERBLOCK	
(FIX) . . . . .	A-24
BAD CYLINDER GROUPS (FIX) . . . . .	A-25
Phase 6: Salvage Cylinder Groups . . . . .	A-25
Cleanup . . . . .	A-25
"f" files, "c" icont, "b" used, "r" free ("y" frags, "z"	
blocks) . . . . .	A-25
***** MARKING FILE SYSTEM CLEAN ***** . . . . .	A-25
***** FILE SYSTEM WAS MODIFIED ***** . . . . .	A-25
***** REBOOT HP-UX ***** . . . . .	A-26



## Using the fsck Command

---

### Introduction

The file system consistency check (`/etc/fsck`) corrects inconsistent information in your file system.

You must have a thorough understanding of the file system before making any `fsck` decisions. Read “File System Implementation” in the “Concepts” chapter before proceeding. Do not execute this command from a diskless node if you are on an HP-UX cluster.

`fsck` should be performed in the following situations:

- `fsck` will run automatically during bootup if you had an unclean shutdown. This will be the case if you had a system crash or power failure.

The `fsck` will run as the system enters run-level 2. As shipped, your system will do this automatically if it detects an improper shutdown, via the `bcheckrc` entry in `/etc/inittab`. An improper shutdown means you didn't shut down your system using the `shutdown` or `reboot` commands described in “Shutting Down the System” in Chapter 3. `fsclean` is used to detect an improper shutdown.

- You should run `fsck` *any* time you suspect problems with the HP-UX file system.

There is no specific problem you can relate to file system corruption. If your HP-UX system doesn't behave like you think it should, run `fsck`.

- You should run `fsck` occasionally even though `fsckclean` (run from `/etc/bcheckrc`) does not indicate the need to do so (particularly if your system exhibits unusual behavior).

The backup script has an option to run `fsck -n` after the backup is performed; you should use this option. `fsck -n` reports discrepancies, but does not fix them, so it can be safely run by `cron`. Note, though, that if `fsck` is run by `cron`, the consistency checking is meaningful only if the file system is otherwise quiescent while `fsck` is running.

`fsck` should be executed using an *unmounted character special device file*, not a block special device file, except when checking the root file system. If you shutdown the system using `shutdown`, all file systems except root will be unmounted and you will be in run-level `s`. Refer to “Adding Peripheral Devices” in Chapter 4 for a discussion on block and character devices, and naming conventions for device files.

The `fsck` program, when run on the root file system, *MUST use the block device* (that is, `dev/dsk/0s0`). If the character device is used, `fsck` doesn’t know that it was the root file system, and won’t issue a reboot instruction when needed.

Unless you explicitly name a file system to be checked, `fsck` will check all the `hfs` entries listed in the `/etc/checklist` file.

Only the system administrator should run `fsck`. If this check discovers an inconsistency, corrective action must be taken. The `fsck` program can be run with any of several options:

default            Interactive mode.

The interactive mode allows you to choose whether or not to perform a corrective action.

-p Preening option (lower-case p).

This option fixes many potential problems, but never removes data. When you preen the system, you are not running interactively. `fsck` will decide what to do, and if it can't deal with a situation, it will terminate. For the inconsistencies the preening option can fix, it will print a message identifying the file system, and the corrective action taken. The preening option will fix the following inconsistencies:

- unreferenced inodes
- unreferenced files and fifos
- link counts in inodes too large
- missing blocks in the free list
- blocks in the free list also in files
- wrong counts in the superblock
- clean byte marked wrong

Other problems will cause `fsck -p` to terminate and prompt for manual execution of `fsck`. Since it was created to be run non-interactively, when it finishes it does not tell you if you need to reboot the system. You must determine that from the return code (refer to *fsck(1M)* in the *HP-UX Reference*).

-P This preening option (upper-case P) is the option used if `bcheckrc` detects an unclean shutdown during system boot. This option operates the same as the `preen (-p)` option except all file systems which were cleanly unmounted will not be checked. For a discussion on clean file systems refer to the chapter "System Startup and Shutdown", the section, "HP-UX Starts the Init Process". Using this option can greatly decrease the amount of time required to reboot a system which was brought down cleanly.

-y "Yes" option.

Using the `-y` option can be very dangerous. This option causes `fsck` to answer YES to all questions, which might remove data. *Do not use the `-y` option if you have important data on your file system unless you have first used the `-n` option and understand the potential damage.*

**-n** “No” option.

Using the **-n** option causes **fsck** to answer **NO** to all questions. This will never remove data, so is very safe. You can use the **-n** option any time; multi-user, single-user, or background.

If you use **fsck** with the **-n** option in multi-user mode, **fsck** may report some inconsistencies due to file system action. However, you will never damage your system.

**-q** Quiet option.

**fsck** prints only the messages that require a response.

**-b** Use an alternate superblock.

Use this option when the superblock is corrupted. A list of alternate superblocks for the file system is in the file `/etc/sbtab`.

When executing **fsck** to perform repairs, the system should *always* be in the system administrative state (run-level **s**) and quiescent (inactive and not being written on). Running **fsck** when there is file system activity may cause loss of data. The **shutdown** command will leave the system in run-level **s**.

A directory with the name `/lost+found` must exist on the file system being examined *before* **fsck** is run. `/lost+found` should have been created on your root volume when you installed HP-UX on your system. The **fsck** command uses this directory for any problem files that it finds. After you run **fsck**, examine the files placed in `/lost+found` and move them to where they belong or remove them. You should clear the `/lost+found` directory before you execute **fsck** again.

To place these files, follow this procedure:

1. Change to the `/lost+found` directory for the file system. For example, `cd /lost+found` for the root file system.
2. Find out what type of file it is (executable, text, etc), and who owns the file, by typing:

```
file *  
ll *
```

If the file is text, see what's in it by typing:

```
more filename
```

3. If the file is executable, then:
  - a. If the file has an RCS ID string, the `what` command will list it.
  - b. If the file does not have an RCS ID string, use the `strings` command to print the literal strings from the file. The strings (for example, error message strings) may help identify the owner.
4. From this information, determine where the file belongs, or who it belongs to, and move the file to the correct directory.

---

## fsck Errors That May Occur

`fsck` is a multi-pass *File System ChecK* program. Each phase of the `fsck` program invokes a different file system pass. After the initial setup, `fsck` performs successive phases over each file system, checking blocks and sizes, pathnames, connectivity, reference counts, and the free-block map (possibly rebuilding it), and doing some cleanup.

Refer to the “Introduction” section of this appendix for a discussion on the different options for `fsck`.

When an inconsistency is detected while running interactively, `fsck` reports the error condition. If a response is required, `fsck` prints a prompt message and waits for a response. When preening, `fsck` will choose a response and note it on the screen. In this section, each error message and possible responses are presented.

The error conditions are organized by the phase of the `fsck` program in which they can occur. The error conditions that may occur in more than one phase are discussed in “Initialization Phase Errors” below.

## **Initialization Phase Errors**

During the initialization phase, before the file system check can be performed, tables have to be set up and certain files opened. This section lists error conditions resulting from command line options, memory requests, opening of files, status of files, file system size checks, and creation of the scratch file. All of the initialization errors are fatal if you are preening. Refer to the *fsck(1M)* entry in the *HP-UX Reference* for further information.

### **“C” option?**

The character represented by `C` is not a legal option to `fsck`. Legal options are `-b`, `-y`, `-n`, `-q`, and `-p`. `fsck` terminates if this error condition occurs. Refer to the *fsck(1M)* entry in the *HP-UX Reference* for further information.

### **cannot alloc NNN bytes for “XXX”**

`XXX` is either `blockmap`, `freemap`, `statemap`, or `lncntp`. `fsck`'s request for memory failed. This should never happen. `fsck` terminates on this error condition. Contact your local HP Sales and Service Office for assistance.

### **Can't open checklist file:**

The default file system checklist file (`/etc/checklist`) cannot be opened for reading. `fsck` terminates on this error condition. Check for the existence of the file, and the access modes of the file.

### **Can't stat root**

`fsck`'s request for statistics about the `root` directory (`/`) failed. This should never happen. `fsck` terminates on this error condition. Contact you local HP Sales and Service Office for assistance.



### **Can't stat ...**

### **Can't make sense out of ...**

`fsck`'s request for statistics about the file system failed. When running manually, it ignores this file system and continues checking the next file system given. If this happens, check for the existence, and the access modes, of the file system.

### **“file\_system\_name” is not a block or character device; OK**

You have given `fsck` a regular file name by mistake. You should check the file type of the file system. Possible responses to the OK prompt are:

- |     |   |
|-----|---|
| YES | Ignore this error condition.  |
| NO  | Ignore this file system and continue checking the next file system given. |

### **Can't open ...**

The file system listed cannot be opened for reading. When running manually, it ignores this file system and continues checking the next file system given. Check the access modes of the file system.

### **“file\_system\_name”: (NO WRITE)**

Either the `-n` flag was specified, or `fsck` failed in its attempt to open the file system “file\_system\_name” for writing. When running manually, all the diagnostics are printed, but no modifications are attempted to fix them.

## Other messages:

MAGIC NUMBER WRONG  
NCG OUT OF RANGE  
CPG OUT OF RANGE  
NCYL DOES NOT GIVE WITH NCG\*CPG  
SIZE PREPOSTEROUSLY LARGE  
TRASHED VALUES IN SUPER BLOCK

and will be followed by the message:

*file\_system*: BAD SUPER BLOCK: *superblock\_address*  
USE -b OPTION TO FSCK TO SPECIFY LOCATION OF AN ALTERNATE  
SUPER-BLOCK TO SUPPLY NEEDED INFORMATION; SEE *fsck(1M)*.

The superblock has been corrupted. An alternative superblock must be used. Refer to the discussion on alternative superblocks under the “Superblock Consistency” section in Chapter 2. Alternate superblock locations are recorded in the file */etc/sbtab*.

## INTERNAL INCONSISTENCY: “message”

An internal problem occurred in *fsck*. “*message*” will indicate the problem. This should never happen. Contact your local HP Sales and Service Office for assistance.

## CAN NOT SEEK: BLK “bn” (CONTINUE)

*fsck*’s request for moving to the specified block number in the file system failed. This should never happen. Contact your local HP Sales and Service Office for further assistance.

Possible responses to the CONTINUE prompt are:

- |     |   |
|-----|---|
| YES | Attempt to continue to run the file system check. Often, however, the problem will persist. This error condition will not allow a complete check of the file system. A second run of <i>fsck</i> should be made to re-check this file system. |
| NO  | Terminate the program.  |

### **CAN NOT READ: BLK ... (CONTINUE)**

`fsck`'s attempt to read a specified block number in the file system failed. This can happen if you have previously interrupted `fsck` before it finished making repairs. Contact your local HP Sales and Service Office for further assistance.

Possible responses to the CONTINUE prompt are:

- YES                    Attempt to continue to run the file system check. Often, however, the problem will persist. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system.
- NO                    Terminate the program.

### **CAN NOT WRITE: BLK ... (CONTINUE)**

`fsck`'s attempt to write a specified block number in the file system failed. The disc is probably physically write-protected. Remove write protections from the disc, and re-run `fsck`.

Possible responses to the CONTINUE prompt are:

- YES                    Attempt to continue to run the file system check. Often, however, the problem will persist. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system.
- NO                    Terminate the program.

### **Phase 1 Errors: Check Blocks and Sizes**

This phase concerns itself with the inode list. This section lists error conditions resulting from checking inode types, setting up the zero-link-count table, examining inode block numbers for bad or duplicate blocks, checking inode size, checking block count, and checking inode format. All errors in phase 1 are fatal if you are preening the file system, except for `INCORRECT BLOCK COUNT` and `BAD INDIRECT ADDRESS`.

## **CG ... : BAD MAGIC NUMBER**

The magic number of cylinder group is wrong. This usually indicates that the cylinder group maps have been destroyed. When running manually, the cylinder group is marked as needing to be reconstructed.

## **UNKNOWN FILE TYPE I= ... (CLEAR)**

The mode word of the inode indicates that the inode is not a character special, block special, regular, network special, fifo, symbolic link, or directory inode.

Possible responses to the CLEAR prompt are:

- YES            De-allocate the inode by zeroing its contents. This will always invoke the UNALLOCATED error condition in Phase 2 for each directory entry pointing to this inode.
- NO            Ignore this error condition.

## **LINK COUNT TABLE OVERFLOW (CONTINUE)**

An internal table for `fsck` containing allocated inodes with a link count of zero has no more room.

Possible responses to the CONTINUE prompt are:

- YES            Continue with the program. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system. If another allocated inode with a zero link count is found, this error condition is repeated.
- NO            Terminate the program.

## **“block\_number” BAD I= ...**

The inode represented by “I=” contains the block number, “block\_number”. This block number is out of the range of the file system. This error condition may invoke the EXCESSIVE BAD BLKS error condition in phase 1 if this inode has too many block numbers outside the file system range. This error condition will always invoke the BAD/DUP error condition in Phase 2 and Phase 4.

### **EXCESSIVE BAD BLKS I= ... (CONTINUE)**

There are more than 10 blocks with a block number out of the range of the file system associated with the inode.

Possible responses to the CONTINUE prompt are:

- YES            Ignore the rest of the blocks in this inode and continue checking with the next inode in the file system. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system.
- NO            Terminate the program.

### **“block\_number” DUP I= ...**

The inode contains block number, “block\_number”, which is already claimed by another inode. This error condition may invoke the **EXCESSIVE DUP BLKS** error condition in phase 1 if this inode has too many block numbers claimed by other inodes. This error condition will always invoke Phase 1b and the **BAD/DUP** error condition in Phase 2 and Phase 4.

### **EXCESSIVE DUP BLKS I= ... (CONTINUE)**

There are more than 10 blocks claimed by other inodes.

Possible responses to the CONTINUE prompt are:

- YES            Ignore the rest of the blocks in this inode and continue checking with the next inode in the file system. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system.
- NO            Terminate the program.

## **DUP TABLE OVERFLOW (CONTINUE)**

An internal table in `fsck` containing duplicate block numbers is full.

Possible responses to the `CONTINUE` prompt are:

**YES**            Continue with the program. This error condition will not allow a complete check of the file system. A second run of `fsck` should be made to re-check this file system. If another duplicate block is found, this error condition will repeat.

**NO**            Terminate the program.

## **PARTIALLY ALLOCATED INODE I= ... (CLEAR)**

The inode is neither allocated nor unallocated.

Possible responses to the `CLEAR` prompt are:

**YES**            De-allocate the inode by zeroing its contents.

**NO**            Ignore this error condition.

## **INCORRECT BLOCK COUNT I= ... (CORRECT)**

The block count for the inode, `inode_number`, is `X` blocks, but should be `Y` blocks. When you are preening the count is corrected.

Possible responses to the `CORRECT` prompt are:

**YES**            Replace the block count of the inode with `Y`.

**NO**            Ignore this error condition.

## **BAD INDIRECT ADDRESS: IND BLOCK n[m] = val I= ... (CORRECT)**

An indirect address block, allocated in the inode indicated by “I=”, contains block address for regions beyond the allocated size of the file. When you are preening, these entries are zeroed.

Possible responses to the CORRECT prompt are:

- YES               Zero the entry.
- NO                Ignore this error condition. Later attempts by the operating system to extend the file into this region may cause a system crash.

## **Phase 1b: Rescan for More Dups**

When a duplicate block is found in the file system, the file system is rescanned to find the inode which previously claimed that block. This section lists the error condition when the duplicate block is found.

### **“block\_number” DUP I= ...**

The inode contains the block number, “block\_number”, which is already claimed by another inode. This error condition will always invoke the BAD/DUP error condition in Phase 2. You can determine which inodes have overlapping blocks by examining this error condition and the DUP/BAD error condition in Phase 1.

## **Phase 2: Check Path Names**

This phase concerns itself with removing directory entries pointing to error conditioned inodes from Phase 1 and Phase 1b. This section lists error conditions resulting from root inode mode and status, directory inode pointers in range, and directory entries pointing to bad inodes. All errors in this phase are fatal if you are preening your file system.

## **ROOT INODE UNALLOCATED. TERMINATING**

The root inode (inode number 2) has no allocated mode bits. This should never happen. The program will terminate. Contact your local HP Sales and Service Office for further assistance.

## **NAME TOO LONG**

The path name shown is too long. This is usually indicative of loops in the file system name space. This can occur if the superuser has made circular links to directories. The offending links must be removed.

## **ROOT INODE NOT DIRECTORY (FIX)**

The root inode (inode number 2) is not directory inode type.

Possible responses to the FIX prompt are:

YES            Change the root inode's type to be a directory. If the root inode's data blocks are not directory blocks, a very large number of error conditions will be produced.

NO            Terminate the program.

## **DUPS/BAD IN ROOT INODE (CONTINUE)**

Phase 1 or Phase 1b found duplicate blocks or bad blocks in the root inode (inode number 2) for the file system.

Possible responses to the CONTINUE prompt are:

YES            Ignore the DUPS/BAD error condition in the root inode and attempt to continue to run the file system check. If the root inode is not correct, then this may result in a large number of other error conditions.

NO            Terminate the program.

## **I OUT OF RANGE I= ... (REMOVE)**

NAME has an inode number (I), which is greater than the end of the inode list.

Possible responses to the REMOVE prompt are:

YES            The directory entry (NAME) is removed.

NO            Ignore this error condition.



### **UNALLOCATED I= ... (REMOVE)**

There are two possible error messages that start like this. One is for directory entries and one is for files. “DIR” or “FILE” has a directory inode (I) without allocated mode bits. The owner, mode, size, modify time, and directory/file name are printed.

Possible responses to the REMOVE prompt are:

YES                   The directory entry is removed.

NO                    Ignore this error condition.

### **DUP/BAD I= ... (REMOVE)**

There are two possible error messages that start like this. One is for directory entries and one is for files. Phase 1 or Phase 1b found duplicate blocks or bad blocks associated with the “DIR” or “FILE” having directory inode, “I”. The owner, mode, size, modify time, and directory are printed.

Generally, the inode with the earliest modify time is incorrect, and should be cleared. To be safe, you should exit `fsck` (`Break`) and inspect the two files to determine which is corrupt, then re-run `fsck`.

Possible responses to the REMOVE prompt are:

YES                   The directory entry is removed.

NO                    Ignore this error condition.

### **ZERO LENGTH DIRECTORY I= ... (REMOVE)**

The directory entry’s size is zero. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the REMOVE prompt are:

YES                   The directory entry is removed. This will always invoke the BAD/DUP error condition in phase 4.

NO                    Ignore this error condition.

### **DIRECTORY TOO SHORT I= ... (FIX)**

The directory entry's size is less than the minimum size for a directory. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

- YES            Increase the size of the directory to the minimum directory size.
- NO            Ignore this error condition.

### **DIRECTORY CORRUPTED I= ... (FIX)**

A directory entry has been found with an inconsistent internal state. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

- YES            Throw away all entries up to the next directory boundary. This drastic action can throw away directory entries, and should be taken only after other recovery efforts have failed.
- NO            Skip to the next directory boundary and resume reading, but do not modify the directory.

### **BAD INODE NUMBER FOR '.' I= ... (FIX)**

The directory entry doesn't have an inode number for '.' which is equal to the inode number. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

- YES            Change the inode number for '.' to be equal to the inode number given after I=.
- NO            Leave the inode number for '.' unchanged

### **MISSING '.' I= ... (FIX)**

The directory doesn't have its first directory entry allocated. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

YES                    Make an entry for '.' with inode number equal to the inode number given after I=.

NO                    Leave the directory unchanged.

### **MISSING '.' I= ...**

#### **CANNOT FIX, FIRST ENTRY IN DIRECTORY CONTAINS ...**

The directory has, as its first entry, the file name given. `fsck` cannot resolve this problem. Exit the `fsck` program, mount the file system (you can force a mount by using `mount -f`), find the file name, and move it to a different directory. Then unmount the file system and run `fsck` again. The owner, mode, size, modify time, and directory name are printed.

### **MISSING '.' I= ...**

#### **CANNOT FIX, INSUFFICIENT SPACE TO ADD '.'**

The directory does not have '.' as its first entry. `fsck` cannot resolve this problem. If this happens, contact your local HP Sales and Service office. The owner, mode, size, modify time, and directory name are printed.

### **EXTRA '.' ENTRY I= ... (FIX)**

The directory has more than one '.' entry. The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

YES                    Remove the extra '.' entry.

NO                    Leave the directory unchanged.

### **BAD INODE NUMBER FOR '..' I= ... (FIX)**

The directory's inode number for '..' does not equal the parent of the inode number (I). The owner, mode, size, modify time, and directory name are printed.

Possible responses to the FIX prompt are:

YES            Change the inode number for '..' to be equal to the parent of the inode given (I).

NO            Leave the inode number for '..' unchanged.

### **MISSING '..' I= ... (FIX)**

The directory doesn't have its second directory entry allocated.

Possible responses to the FIX prompt are:

YES            Make an entry for '..' with inode number equal to the parent of the inode number given (I).

NO            Leave the directory unchanged.

### **MISSING '..' I= ...**

#### **CANNOT FIX, SECOND ENTRY IN DIRECTORY CONTAINS ...**

The directory has, as its second entry, the file name given. `fsck` cannot resolve this problem. The file system should be mounted and the offending entry moved elsewhere. To do this, exit the `fsck` program, mount the file system (you can force a mount by using `mount -f`, find the file name, and move to a different directory. The file system should then be unmounted and `fsck` should be run again.

### **MISSING '..' I= ...**

#### **CANNOT FIX, INSUFFICIENT SPACE TO ADD '..'**

The directory does not have '..' as its second entry. `fsck` cannot resolve this problem. If this happens, contact your local HP Sales and Service office.

### **EXTRA '..' ENTRY I= ... (FIX)**

The directory has more than one entry for '..'.

Possible responses to the FIX prompt are:

- YES                Remove the extra entry for '..'.
- NO                Leave the directory unchanged.

### **Phase 3: Check Connectivity**

This phase concerns itself with the directory connectivity seen in Phase 2. This section lists error conditions resulting from unreferenced directories, and missing or full `/lost+found` directories.

#### **UNREF DIR I= ... (RECONNECT)**

The directory inode (I) was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of the directory inode are printed. If you are preening, the directory is reconnected if its size is non-zero, otherwise it is cleared.

Possible responses to the RECONNECT prompt are:

- YES                Reconnect the directory inode to the file system in the directory for lost files (`/lost+found`). This may invoke the `lost+found` error condition if there are problems connecting the directory inode to `/lost+found`. This may also invoke the `CONNECTED` error condition in phase 3 if the link was successful.
- NO                Ignore this error condition. This error will always invoke the `UNREF` error condition in Phase 4.

#### **SORRY, NO lost+found DIRECTORY**

There is no `/lost+found` directory in the root directory of the file system. `fsck` ignores the request to link a directory in `/lost+found`. This will always invoke the `UNREF` error condition in Phase 4. Check access modes of `/lost+found`. This error is fatal if you are preening the system.

## **SORRY. NO SPACE IN `lost+found` DIRECTORY**

There is no space to add another entry to the `/lost+found` directory in the root directory of the file system. `fsck` ignores the request to link a directory in `/lost+found`. This will always invoke the UNREF error condition in Phase 4. Clean out unnecessary entries in `/lost+found` or make `/lost+found` larger and try again. This error is fatal if you are preening the system.

## **DIR I= ... CONNECTED. PARENT WAS I= ...**

This is an advisory message indicating a directory inode was successfully connected to the `/lost+found` directory. The parent inode of the directory inode is replaced by the inode number of the `/lost+found` directory.

## **BAD CONTINUATION INODE NUMBER. I= ... CLEAR?**

If no optional ACL entries are on your system, it is an inode created with the “proto” option in `mkfs`. If you have optional ACL entries on your system, you may be losing that ACL data. This field was unused previous to 6.5, but is required for C2 security.

**YES** Clear the continuation inode field. You must clear it or you may receive a kernel panic later, but you may want to write down the `I#` to check if there is any ACL information in it. If you just removed it with `fsck` there will not be any ACL information. If you do not have any ACL entries on your system, simply answer YES and ignore the I number.

**NO** Does not clear it. This may later cause a kernel panic.

## **Phase 4: Check Reference Counts**

This phase concerns itself with the link count information seen in Phase 2 and Phase 3. This section lists error conditions resulting from unreferenced files, missing or full `/lost+found` directory, incorrect link counts for files, directories, or special files, unreferenced files and directories, bad and duplicate blocks in files and directories, and incorrect total free-inode counts. All errors in this phase are correctable if you are preening your file system, except if you run out of space in `/lost+found`.

## **UNREF CONTINUATION INODE CLEAR?**

Yes is the only correct response. If you had a bad continuation inode number in pass 3 you probably have lost ACL data.

## **BAD LINK COUNT IN CONTINUATION INODE. I= ... (LINK COUNT SHOULD BE 1) CORRECT?**

The *only* correct response is YES. If using the preen option, this is a warning, and will automatically be corrected.

## **UNREF FILE I= ... (RECONNECT)**

The inode (I) was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of the inode are printed. If you are preening, the file is cleared if either its size, or its link count, is zero. Otherwise it is reconnected.

Possible responses to the RECONNECT prompt are:

- YES                Reconnect the inode to the file system in the directory for lost files (usually /lost+found). This may invoke the lost+found error condition if there are problems connecting the inode to /lost+found.
- NO                Ignore this error condition. This will always invoke the CLEAR error condition in phase 4.

## **(CLEAR)**

The inode mentioned in the immediately previous error condition cannot be reconnected. If you are preening, this error cannot occur, since lack of space to reconnect files is a fatal error.

Possible responses to the CLEAR prompt are:

- YES                De-allocate the inode mentioned in the immediately previous error condition by zeroing its contents.
- NO                Ignore this error condition.

## **SORRY. NO lost+found DIRECTORY**

There is no `/lost+found` directory in the `root` directory of the file system. `fsck` ignores the request to link a file in `/lost+found`. This will always invoke the `CLEAR` error condition in phase 4. Check access modes of `/lost+found`. If you are preening your file system, this error is fatal.

## **SORRY. NO SPACE IN lost+found DIRECTORY**

There is no space to add another entry to the `/lost+found` directory in the `root` directory of the file system. `fsck` ignores the request to link a file in `/lost+found`. This will always invoke the `CLEAR` error condition in phase 4. Check size and contents of `/lost+found`. This error is fatal if you are preening your file system.

## **LINK COUNT ... (ADJUST)**

The link count for the file, directory, or inode is one link count (`COUNT=`) but should be a different link count (`SHOULD BE`). The owner, mode, size, and modify time are printed. If you are preening, the link count is adjusted.

Possible responses to the `ADJUST` prompt are:

- YES            Replace the link count of the file in the inode with the “`SHOULD BE`”: number.
- NO            Ignore this error condition.

## **UNREF ... (CLEAR)**

The file or directory with the inode number, “`I`”, was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of the inode are printed. If you are preening, the inode is cleared, since this is a file that was not connected because its size or link count was zero.

Possible responses to the `CLEAR` prompt are:

- YES            De-allocate inode by zeroing its contents.
- NO            Ignore this error condition.



## **BAD/DUP ... (CLEAR)**

Phase 1 or Phase 1b found duplicate blocks or bad blocks associated with the file or directory inode given in "I=". The owner, mode, size, and modify time of the inode are printed. This error will not occur if you are preening, since it would have caused a fatal error earlier.

Possible responses to the CLEAR prompt are:

YES                De-allocate the inode by zeroing its contents.

NO                 Ignore this error condition.

Often deleting only one of the files containing DUPS will cure the problem. `fsck` should be re-run to confirm that the problem was fixed. A NO means that `fsck` must be re-run to finish cleaning up the file system.

## **FREE INODE COUNT WRONG IN SUPERBLK (FIX)**

The actual count of the free inodes does not match the count in the super-block of the file system. If you are preening, the count is fixed.

Possible responses to the FIX prompt are:

YES                Replace the count in the super-block by the actual count.

NO                 Ignore this error condition.

## **Phase 5: Check Cylinder Groups**

This phase concerns itself with the free-block maps. This section lists error conditions resulting from allocated blocks in the free-block maps, free-blocks missing from free-block maps, and the total free-block count not matching the count contained in the super-block.

### **CG ... : BAD MAGIC NUMBER**

The magic number of the cylinder group is wrong. This usually indicates that the cylinder group maps have been destroyed. When running manually, the cylinder group is marked as needing to be reconstructed. If you are preening your system, this error is fatal.

## **EXCESSIVE BAD BLKS IN BIT MAPS (CONTINUE)**

You should never get this message. If you do, contact your local HP Sales and Service office.

### **SUMMARY INFORMATION “t” BAD**

where *t* is one or more of:

(INODE FREE)

(BLOCK OFFSETS)

(FRAG SUMMARIES)

(SUPER BLOCK SUMMARIES)

The indicated summary information was found to be incorrect. This error condition will always invoke the BAD CYLINDER GROUPS condition in phase 5. If you are preening, the summary information is recomputed.

### **“x” BLK(S) MISSING**

A number of blocks (x) that are unused by the file system were not found in the free-block maps. This error condition will always invoke the BAD CYLINDER GROUPS condition in phase 5. If you are preening, the block maps are rebuilt.

### **FREE BLK COUNT WRONG IN SUPERBLOCK (FIX)**

The actual count of free blocks does not match the count in the super-block of the file system. If you are preening, the counts are fixed.

Possible responses to the FIX prompt are:

YES                    Replace the count in the super-block by the actual count.

NO                     Ignore this error condition.

## **BAD CYLINDER GROUPS (FIX)**

Phase 5 has found bad blocks in the free-block maps, duplicate blocks in the free-block maps, or blocks missing from the file system. If you are preening, the cylinder groups are reconstructed.

Possible responses to the FIX prompt are:

YES                Replace the actual free-block maps with new free-block maps.

NO                 Ignore this error condition.

## **Phase 6: Salvage Cylinder Groups**

This phase concerns itself with reconstructing the free-block maps. No error messages are produced.

## **Cleanup**

Once a file system has been checked, a few cleanup functions are performed. This section lists advisory messages about the file system and modify status of the file system.

### **“f” files, “c” icont, “b” used, “r” free (“y” frags, “z” blocks)**

This message indicates that the file system just checked has a total of “f” files, using “b” fragment-sized blocks, with “r” fragment-sized blocks available (free) for use. It also has a total of “c” continuation inodes (which also is the number of optional ACL entries). The numbers in parenthesis divides the free count into “y” free fragments and “z” free full-sized blocks.

No action is required on your part.

### **\*\*\*\*\* MARKING FILE SYSTEM CLEAN \*\*\*\*\***

Indicates that the fil system clean byte is being set to FS\_CLEAN.

### **\*\*\*\*\* FILE SYSTEM WAS MODIFIED \*\*\*\*\***

This message indicates that the current file system was modified by `fsck`. If this file system is mounted or is the current root file system, `fsck` should be halted

and HP-UX re-booted. If HP-UX is not re-booted immediately, the work done by `fsck` may be undone by the in-core (memory) copies of tables HP-UX keeps.

**\*\*\*\*\* REBOOT HP-UX \*\*\*\*\***

This message indicates that the root file system has been modified by `fsck`. If HP-UX is not re-booted immediately, the work done by `fsck` may be undone by the in-core (memory) copies of tables HP-UX keeps. If you are preening, `fsck` will exit with a code of 4. The `bcheckrc` script interprets an exit code of 4 by executing the reboot command.

---

**Note**            If you are preening, you will not get this message. `bcheckrc` will correctly interpret the exit code and automatically reboot, but if you execute `fsck -p` outside the `bcheckrc` program, you must check the return code to see if the system needs to be rebooted.

---

# Contents

---

## B. Troubleshooting

System Run-Time Errors . . . . .	B-2
Kernel Errors . . . . .	B-2
“panic: <message>” . . . . .	B-2
“pid <num>: in xalloc: killed, not enough swap space.” “sorry, pid <num> was killed, not enough swap space.” . . . . .	B-3
“<name>: table is full” . . . . .	B-3
“no more processes” . . . . .	B-3
“danger: mfree map overflow” . . . . .	B-3
File System Errors . . . . .	B-4
Error Messages . . . . .	B-5
Booting Diagnostic Errors . . . . .	B-5
“Not enough memory” “Panic: not enough memory to powerup” . . . . .	B-5
“WARNING: avail mem too low, may deadlock” . . . . .	B-5
When you boot, you do not receive a login prompt . . . . .	B-6
Error Messages During Login . . . . .	B-6
“No Shell” . . . . .	B-6
System Will Not Accept a Password . . . . .	B-7
Unable to Access a Peripheral Device . . . . .	B-8
Dealing with an Unresponsive Terminal . . . . .	B-9
Backup Recovery Errors . . . . .	B-10
“cpio: Out of phase.” “Shall I try to resynch (y,n)?:?” . . . . .	B-10
Disk Hard Error . . . . .	B-11
“<command>: not found” (in Bourne shell) “<command>: Command not found” (in C shell) “ksh: <command>: not found” (in Korn shell) . . . . .	B-11
“Unable to access ttyxx” . . . . .	B-13

Update Errors . . . . .	B-13
Errors During Update . . . . .	B-13
“Invalid Destination directory.” . . . .	B-13
“Source special file <source device name> does not exist.” . . . .	B-13
“Invalid parameters for source device.” . . . .	B-13
“Cannot read table of contents on <source device name>” . . . . .	B-13
“Cannot read <fileset> on <source device name>” . . . . .	B-14
“Choice not valid” . . . . .	B-14
“Copy of <filename> failed.” . . . .	B-14
“FATAL ERROR: Too many processes running for single-user state.” “Execute the shutdown command before performing the Update.” . . . .	B-14
“This product may not be loaded using the -n option of update.” “Please exit and rerun update without the -n option” . . . . .	B-14
“Permission denied.” . . . .	B-14
“WARNING: Update can’t check for single-user state.” . . . .	B-15
“cdfcustomize:” “Warning: <filename> should be a CDF rather than an ordinary file” . . . . .	B-15
To check your tapes: . . . . .	B-16
HP-UX Cluster Troubleshooting . . . . .	B-17
Unsolved Problems . . . . .	B-17
Flowchart Format . . . . .	B-18
Diskless Boot Troubleshooting . . . . .	B-18
Initial Steps to Troubleshooting Diskless Boot Problem (Flowchart A) . . . . .	B-20
(Flowchart A Cont.) . . . . .	B-21
/etc/rbootd Not Running (Flowchart B) . . . . .	B-23
(Flowchart B Cont.) . . . . .	B-24
/etc/rbootd Not Running (Flowchart C) . . . . .	B-26
(Flowchart C Cont.) . . . . .	B-27
Check for Multiple LAN Interfaces in Server (Flowchart D) . . . . .	B-29
(Flowchart D Cont.) . . . . .	B-30
/etc/clusterconf Troubleshooting (Flowchart E) . . . . .	B-32
(Flowchart E Cont.) . . . . .	B-33

Run Hardware Diagnostics (Flowchart F) . . . . .	B-35
(Flowchart F Cont.) . . . . .	B-36
Other Possible Problems (Flowchart G) . . . . .	B-38
(Flowchart G Cont.) . . . . .	B-39
Some Additional Hardware Checks (Flowchart H) . . . . .	B-41
(Flowchart H Cont.) . . . . .	B-42
Possible Boot Daemon Error Messages . . . . .	B-43
(/etc/rbootd Error Messages Cont.) . . . . .	B-44
(/etc/rbootd Error Messages Cont.) . . . . .	B-45
Possible Panic, Retry, and Other System Messages on	
Diskless Cnodes . . . . .	B-46
(System Messages Cont.) . . . . .	B-47
(System Messages Cont.) . . . . .	B-48
(System Messages Cont.) . . . . .	B-49
(System Messages Cont.) . . . . .	B-50
(System Messages Cont.) . . . . .	B-51
System Messages Cont.) . . . . .	B-52
LAN Cable Break Detection Behavior . . . . .	B-52
Unable to Login or Rlogin to a Diskless Cnode . . . . .	B-55





# B

## Troubleshooting

---

This appendix gives guidelines for what to do when specific error messages or conditions arise on your system. Actual error messages are shown inside quotes.<sup>1</sup> If a part of the error message is shown inside brackets, <>, that part of the message changes with the specific circumstance. Error conditions are not enclosed in quotes.

---

<sup>1</sup> Error messages in tables are in computer font.

---

## System Run-Time Errors

These are errors that you might see on the console just before a system crash, or problems you might encounter due to a corrupt file system.

### Kernel Errors

The following are messages and conditions you may see if there is a problem with your kernel.

#### “panic: <message>”

This means that something is wrong in the kernel or hardware.

One possible reason for the following message is a driver I/O error:

```
panic: bread: size 0
```

Try to do an `fsck`. If `fsck` terminates with: `CAN NOT READ: ... (CONTINUE)`, you should suspect a disk hardware problem.

If the system runs out of callouts (refer to the `nccallout` entry in Appendix D), the following fatal error will be printed to the console:

```
panic: timeout table overflow
```

If, in addition to the “panic” message, you get a system dump on your screen, you should copy all the information on your screen dump onto paper. This will help HP support personnel figure out what the problem is.

If you see one of the following messages, you most likely have bad blocks on your disk:

```
panic: IO err in push  
or  
PID: killed on swap error
```

Back up your file system, re-initialize your disk (`mediainit`) to spare out the bad blocks, re-create a file system (`newfs`), and restore your files from the backup.

If you see the following message:

```
panic: process_name: continuation inode is wrong file type; fsck the disc
```

This can happen because the 6.5 kernel uses a previously unused inode. The `fsck` program will fix the problem. You must manually `fsck` the disk on reboot, and must answer YES to the `fsck` message:

```
BAD CONTINUATION INODE NUMBER.  I= ...  
CLEAR?
```

**“pid <num>: in xalloc: killed, not enough swap space.”**

**“sorry, pid <num> was killed, not enough swap space.”**

This means there is not enough swap space to run your process. Your swap space must be as big as the combined total of all the applications running. Refer to Chapter 6 for information on configuring more swap space.

**“<name>: table is full”**

One of the kernel tables, whose size is configurable, is full. This message is printed on the console for information only. If it happens frequently you should probably configure a kernel with a larger table. Refer to the <name> entry in Appendix D for more information.

**“no more processes”**

Two possible times to see the above error message are:

- When the total number of processes for a given user is larger than `maxuprc`, this message will be sent to the user `tty`. Refer to the `maxuprc` entry in Appendix D for more information.
- When the total number of processes in the system is larger than `nproc`, the above message will be sent to the `tty` of the user who tries to start the last process(es). Refer to the `nproc` entry in Appendix D for more information.

**“danger: mfree map overflow”**

You will see this message if semaphore usage is heavy and a request for a semaphore set cannot be accommodated. Refer to the `semmap` entry in Appendix D for more information.

## **File System Errors**

If you suspect problems with your file system, run `fsck`. This command is documented in Appendix A. You may also wish to read the section “HFS File System” in the chapter “System Management Concepts”.

---

## Error Messages

This section covers error messages or error conditions you might encounter.

### Booting Diagnostic Errors

Errors from the boot ROM are documented in your computer's hardware manual.

#### **“Not enough memory”**

#### **“Panic: not enough memory to powerup”**

If either of these messages appears when you try to boot the system you must either make more memory available to your system or boot a smaller system (try booting `/SYSBCKUP`). You can make more memory available to your system by:

- correcting any improper memory configuration (refer to the “Memory Configuration Wheel” that came with your computer)
- replacing defective memory cards
- adding more memory cards

After making any necessary changes, try to reboot the system.

You should also check to see if the boot ROM reports the same amount of memory you have installed in the system: if not, you probably have a hardware problem.

#### **“WARNING: avail mem too low, may deadlock”**

If this message appears and you do not receive a login prompt your system needs additional available memory. Refer to the information in the previous section “Not enough memory”.

If this message appears and you are able to login, your system should run. However, your system may have poor performance unless you make more memory available.

## When you boot, you do not receive a login prompt

If, when you boot the system, you do not receive a login prompt, you have come up in the System Administration Mode (run-level s). There are several possibilities:

- no `/etc/passwd` file

Refer to the information under the error message “No Shell”.

- no `/etc/inittab` file

Similar procedure as documented for `/etc/passwd` recovery under “No Shell”. If you do not have a backup copy of `/etc/inittab`, you will find a copy on your recovery system and a copy on your installation media.

- error in the `/etc/inittab` file

A syntax error in `/etc/inittab` results in booting into run-level s. See above solution.

## Error Messages During Login

### “No Shell”

If you get this error message while attempting to log in, it means that `exec` failed from `login`. Possible causes of this failure are:

- Missing or corrupted `/etc/passwd` file. If the `/etc/passwd` file is corrupted, your system will not allow anyone to login in the multi-user mode. There will be a system prompt but the system will respond to any entry with “No Shell”.

Any time you modify the `/etc/passwd` file you should first rename the existing file `/etc/opasswd`. This allows you to duplicate the old file if you made a mistake or when the new file is destroyed. You can copy this file and then `sync` the system.

If you cannot login to your system as superuser, then your system will automatically put you in the system administration mode (run-level s). From this point, you can manipulate the `/etc/passwd` file.

- User access incorrect on the shell or on directories. Check the access permissions on `/bin/sh`, the root directory, the `/bin` directory, and the

user's login directory. Table B-1 shows the correct permissions, the correct owner, and the correct parameter to use with `chmod` to get the correct permissions.

**Table B-1. Access Permissions**

File or Directory	Correct Access	Correct Owner/Group	chmod parameters
/bin/sh	-r-xr-xr-t	bin/bin	4555
/	drwxr-xr-x	root/root	555
/bin	drwxr-xr-x	root/other	755
User's home directory	The permissions should be read, write, and execute for the owner. Access permissions for group and others do not matter.	The user's login name.	7xx

### **System Will Not Accept a Password**

If you see the following symptoms, the `/dev/tty` file for the terminal is probably destroyed:

- The login prompt is displayed.
- When a valid login name is entered, the password prompt is displayed, but the system immediately moves the cursor to the next line and you will not be able to enter a password.

To fix this problem, perform the following:

1. Recreate the `/dev/tty` file by entering a `mknod` command. For example:

```
mknod /dev/tty c 2 0
```

2. Set the file permission mode for the `/dev` directory:

```
chmod 755 /dev
```

3. Check the file systems for the correct file permission modes on other system files. Verify that the file permission modes will prevent other users from changing or destroying important system files accidentally. Most system files require the same access permission as set in step 2.

## Unable to Access a Peripheral Device

If you cannot access a peripheral device, such as a printer or a disk, there are several possibilities, both hardware and software. Unfortunately, the problem can come from so many different places that there is no concise list of error messages you might see. The following is a list of areas to check to help pinpoint the source of the problem:

### ■ Hardware Problem

1. Check to make sure the device is turned on and has completed any initialization cycle.
2. Make sure all cables are correctly connected, and that the device is plugged in.
3. If you have ruled out all other possible problems, use your hardware documentation to troubleshoot the device.

- You cannot access the hardware correctly There are many possible scenarios for this. One example is when using the `cpio` command with an HP 9144 1/4 inch tape drive. Because `cpio` defaults to 512-byte blocks, but the tape drive needs 1-Kbyte blocks, you will get the following error:

```
no such device
```

You need to specify the block size in the `cpio` command line.



- The device file is incorrect for the device. Generally you will get the following message:

**Cannot open device**

The most common reasons for this are:

- wrong address
- wrong major or minor number
- wrong block/character mode
- if you are on an HP-UX cluster, you may have the wrong cnode access permission

Remake the device file, checking to see that you use the correct information for major and minor numbers (refer to Chapter 4)

If you are on an HP-UX cluster, type `ll -H` for the device file. Check to see if the cnode id field matches the cnode from which you are trying to access the device. If not, you need to create a device file from the appropriate cnode so you have the correct access permissions. Refer to Chapter 2, the section “Cluster Concepts”, the subsection “Cnode-Specific Device Files” for more information.

- Device driver not configured. Go back to Chapter 4, the section that discusses adding your peripheral, to find out what the correct device driver is. Add it to your kernel using the procedures in Chapter 6, “Configuring Device Drivers and I/O Cards Using reconfig”.

## Dealing with an Unresponsive Terminal

Some common user errors will result in the terminal being left in an unusable configuration. One example of this would be when a user tries to display the contents of a binary file (for example, an executable program) on a terminal.

If, for whatever reason, a terminal will not respond (or does not appear to respond) to your commands, several alternative solutions are available:

- The first is to simply log off the system (using `exit`, `CTRL-D`, or whatever the recommended method for the shell you are running is), and then login again. Generally this will clear up any problems.
- A second solution is to do a hard reset of your terminal. If you do a hard reset, the tabs may be cleared. To determine if you need to reset the tabs,

execute `ls`. If the files are listed in a column instead of across the page, your tabs need to be reset. To reset the tabs, execute the `/usr/bin/tabs` command.

- If you have properly set up the `/etc/ttytype` file, the `tset` command will generally cause the terminal to respond again. To use this command, type the following *exactly* as shown—blanks are significant:

```
CTRL-J stty sane erase "^H" kill "^U" echo CTRL-J
```

This command sets the erase character to `CTRL-H` and the kill character to `CTRL-U`. Following the execution of this command, you may not see anything echoed on the screen. When the screen and keyboard response returns, type:

```
tset Return
```

Your terminal should now exhibit proper behavior. This may not work if the terminal is in 8-bit mode.

- Another option is to execute the `tabs` command.

## Backup Recovery Errors

The backup procedure is documented in the chapter “Periodic System Administrator Tasks”. Refer to this procedure for additional information.

**“cpio: Out of phase.”**

**“Shall I try to resynch (y,n)?:?”**

This message comes from `cpio`. It indicates that `cpio` could not successfully read the header.

This can happen if you attempt to restore files from a multi-media backup, but do not begin with media number 1. If you type `y` and press the `Return` key, `cpio` will search for the next good entry.

Press the `Return` key, `cpio` will search for the next header on the media and continue.

## Disk Hard Error

If one of your backup flexible disks has an error, you can still recover most of the files on the disk by using the `dd` command. The following example assumes the name of your flexible disk's device file is `/dev/rdisk/1s0`.

To find out which file is corrupted, use the `cpio` command to get a list of the files on the disk. Look for the last filename printed before the error messages. The next file in the backup might be corrupted too but you will not see its name:

```
cpio -it /dev/rdisk/1s0
```

The flexible disk with read errors will not allow the backup restoration to continue. If the bad flexible disk is copied to another one, the new flexible disk will look just like the old one, except the bad blocks are now readable. Use the following to copy the flexible disk with read errors:

```
dd if=/dev/rsave of=/usr/tmp/flop conv=noerror
```

Now insert the new flexible disk and type:

```
dd if=/usr/tmp/flop of=/dev/rsave bs=18b
```

The new flexible disk will now have all the data from the old flexible disk with one or more corrupted files. The restore can be restarted using `cpio`. If you get a message saying:

```
cpio: Out of phase.  
Shall I try to resynch (y/n)??
```

Type `y` and press the  key to continue. The `cpio` command will look for the next good entry.

When the restore is complete look for the corrupted files and remove them.

**“<command>: not found” (in Bourne shell)**

**“<command>: Command not found” (in C shell)**

**“ksh: <command>: not found” (in Korn shell)**

If you execute a command, but get one of the above errors, there may be one of several things wrong:

- You may have typed the command incorrectly. For example, check to see if you substituted a `1` (the number one) for an `l` (lower-case L). Also, check to see if you typed the correct path.

- You may not have the command's directory location in your PATH variable. To check, type the following:

```
echo $PATH           (for Bourne shell, Korn shell, or PAM)
echo $path          (for C shell)
```

If the command's directory is not in your PATH variable, you can add it by typing the line in column one of Table B-2. Make sure you type the correct line for the shell you are running. If you wish to permanently add this new PATH to your environment, add it to your local environment file.

**Table B-2. PATH Variable**

Shell	Command Line	Local Environment File
Bourne shell (/bin/sh) and Korn shell (/bin/ksh)	PATH=\$PATH: <i>new_directory</i>	.profile
C shell (/bin/csh)	set path=(\$path <i>new_directory</i> )	.cshrc
PAM (/bin/pam)	PATH=\$PATH: <i>new_directory</i>	.environ

- The command does not exist on your system. This may be because:
  - you do not have the fileset containing the command loaded onto your system. Check the appendix, "Partitions and Filesets", for information on checking this.
  - you (possibly accidentally) removed the command from your system. Refer to the appendix, "Partitions and Filesets", for information on how to load the command onto your system.
  - the command is part of some optional software that you do not have. For example, many commands in full HP-UX do not exist in AXE. You would need to purchase the Programming Environment (PE) to have all the HP-UX commands.

## **“Unable to access ttyxx”**

The two possible problems are:

- The kernel does not have the appropriate driver to support the interface card your terminal (port number xx) is connected to.
- The minor number is incorrect in the device file.

Refer to the chapter “Customizing the HP-UX System,” the section “Terminals and Modems”.

---

## **Update Errors**

The following is a list of some the error messages that a user might see while updating HP-UX or installing an optional product with the `/etc/update` program.

There are also some simple tests one can do on tapes and floppies to find obviously bad media; they will not find all problems but they may point out some problems.

### **Errors During Update**

#### **“Invalid Destination directory.”**

Either the user has specified a non-existent directory or the product being unpacked does not allow any directory except `/` to be the destination.

#### **“Source special file <source device name> does not exist.”**

The device file specified for the source does not exist.

#### **“Invalid parameters for source device.”**

The device file specified for the source exists but is not a character device file.

#### **“Cannot read table of contents on <source device name>”**

Update cannot read the contents file from the source device. Make sure that the source device is correct and that the media is in the device specified by the source

device. If you are using a tape make sure it is done conditioning and is locked in the drive. If you have a usable system lifts the media and make sure that it has either a file called CONTENTS or a file called INDEX.

**“Cannot read <fileset> on <source device name>”**

Update cannot read a fileset off of the source device. Make sure that you have the correct media in the correct device. If possible lifts the media and look for a file whose name is the name of the fileset with an 'A' prepended onto the front of it.

**“Choice not valid”**

User tried to use the  softkey when the highlighted selection was . This is not allowed and does not have any specified action (other than to print out this message).

**“Copy of <filename> failed.”**

Update copies some special files after loading all of the other filesets that were selected or loaded. This message means that one of those copies failed and you probably don't have the file named in the message on your system.

**“FATAL ERROR: Too many processes running for single-user state.”**

**“Execute the shutdown command before performing the Update.”**

The user has not shutdown the system before running update.

**“This product may not be loaded using the -n option of update.”**

**“Please exit and rerun update without the -n option”**

The user ran update with the -n option and the product that they are trying to load does not allow updating with the -n (no reboot) option. (The HP-UX AXE and PE products cannot be loaded with the -n option.)

**“Permission denied.”**

You must be superuser to run update, regardless of the mode of the /etc/update file.

**“WARNING: Update can’t check for single-user state.”**

Update could not determine whether or not the system is in init state s; if you are sure you are in init state s you can continue otherwise you should answer no, do a shutdown and rerun update. (This could be caused by a problem with the wc command or the ps command.)

**“cdfcustomize:”**

**“Warning: <filename> should be a CDF rather than an ordinary file”**

You will see this only if you are on an HP-UX cluster. It means that the <filename> is a file that is generally a CDF. Verify that you want the file to be a regular file, or change it to be a CDF.

## To check your tapes:

Here are some ideas for checking tapes or floppies to see if they may be bad.

To find out if an update tape may be bad try the following:

```
lifls -l /dev/update.src
```

The list of lif files will vary with each product but you should not get an error.

If you suspect a particular fileset on the tape is bad you could do:

```
mkdir /tmp/dummyroot
cd /tmp/dummyroot
lifcp -r /dev/update.src:<lif-file-name> - | cpio -idumvx
```

note: *<lif-file-name>* is the name you see when you do a lifls on the media and corresponds to the fileset name with an 'A' prepended onto the front of it.

If you get any error messages from cpio or lif then the tape or flexible disk is probably bad.

One of the following two command lines should work and print out a table. If neither of these files exist then you have a bad media.

```
lifcp -a /dev/update.src:CONTENTS -
or
lifcp -a /dev/update.src:INDEX -
```

If you have an install tape (and a usable system) you can create a block device file which corresponds to /dev/update.src and mount the install media. You can then look at the files and see if they look reasonable. Also you can fsck the media.



---

## HP-UX Cluster Troubleshooting

If a diskless node is not booting or is experiencing system panics, use this appendix to identify and correct the problem. This section has the following organization:

- Diskless Boot Troubleshooting Flowcharts
- Possible Boot Daemon Error Messages
- Possible Panic, Retry, and Other System Messages on Diskless Cnodes
- LAN Cable Break Detection Behavior
- What to do if You Cannot Login or Rlogin to a Diskless Cnode

If you are not familiar with the operation of diskless workstations or the terminology used in discussing clusters of diskless workstations, refer to the concepts chapter in this manual before continuing with the troubleshooting steps.

### Unsolved Problems

If you do not solve the problem after working through all possible troubleshooting steps, call your HP support representative for assistance. Provide as much information about the problem as possible, including the following:

- The activity you were attempting when the error occurred as well as the events leading to and including the problem.
- The cluster configuration. This should include the hardware configuration of all cnodes involved. For each cnode, be sure to give the SPU type and the amount of installed memory. If any of the cnodes have multiple LAN interfaces, be sure to give this information including the select code settings of all the LAN cards. Also describe how the cluster interacts with other standalone systems and with other clusters in your network.
- The version or update information for all software you are running. The version of HP-UX that you are running may be found from the `uname -a` command.
- The error messages or panic message you received. Record all error messages and numbers that appeared both on the diskless cnode and root server.

- The troubleshooting steps you tried.
- The problems you ruled out and why.

## Flowchart Format

The flowcharts on the following pages each have a corresponding set of labeled explanations. You should refer to the explanatory text for more detail when going through the flowcharts.


- Start of Flowchart #



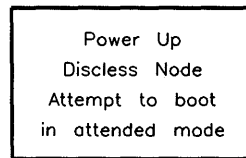
- Go to Flowchart #



- Make a decision



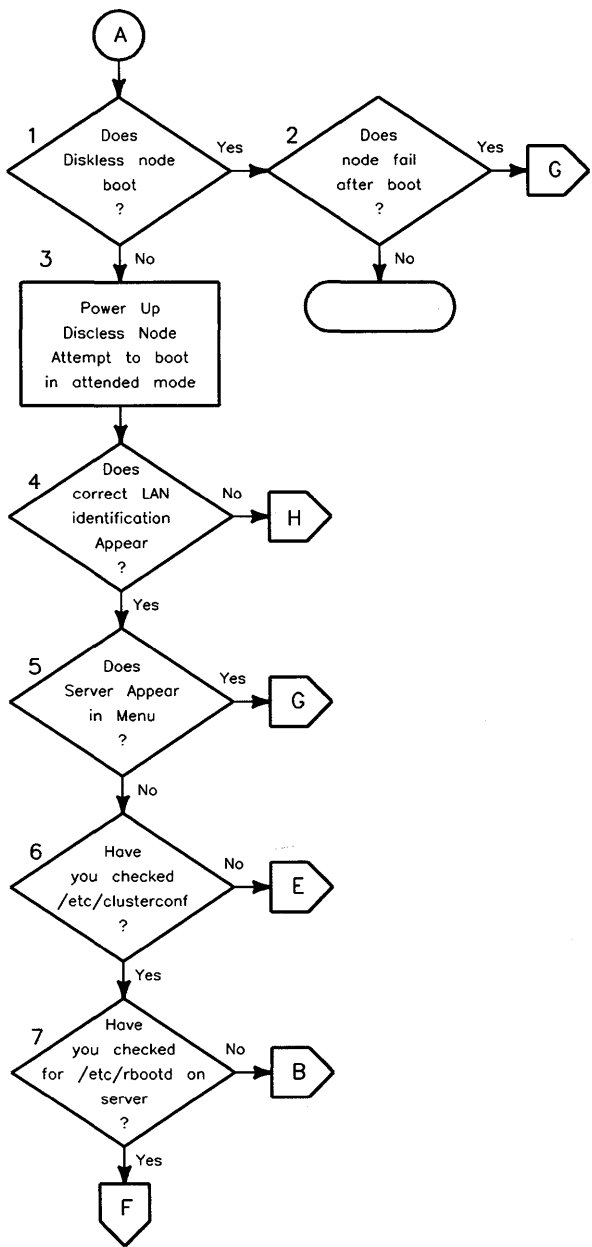
- Perform an action



- Exit Flowchart



## Diskless Boot Troubleshooting



Flowchart A

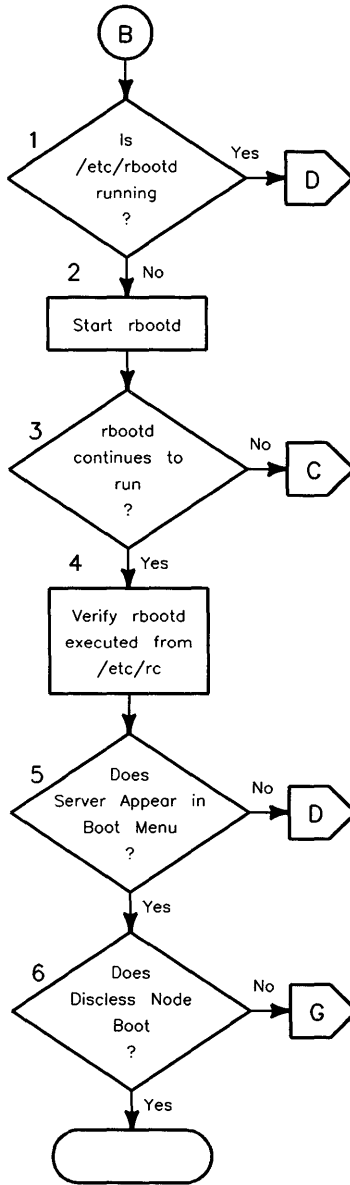
## Initial Steps to Troubleshooting Diskless Boot Problem (Flowchart A)

Begin your troubleshooting efforts with Flowchart A since it helps you determine the best troubleshooting path based on the problem's symptoms.

Question	Yes: Action	No: Action
<p>1. Does the diskless cnode boot successfully?</p> <p>Booting successfully means the diskless cnode successfully loads and executes the operating system, successfully executes the <code>/etc/brc</code>, <code>/etc/bcheckrc</code> and <code>/etc/rc</code> scripts and displays a <code>login:</code> prompt.</p>	See 2.	See 3.
<p>2. Does the diskless cnode boot successfully, but fail later on after the user has logged in?</p>	See Flowchart G.	Since the diskless cnode is not failing, this troubleshooting procedure will not be useful.
<p>3. To begin the troubleshooting process, attempt to boot the diskless cnode in attended mode. This means that you hit the spacebar once the boot messages begin to appear. A menu of bootable systems should be displayed. Refer to the "System Startup and Shutdown" chapter of this manual for more information about attended mode.</p>	See 4.	See 4.
<p>4. Do the boot ROM hardware diagnostics on the left side of the screen show a correct entry for the HP 98643A interface? A correct entry will look something like this:</p> <p style="padding-left: 40px;"><code>HP98643 at 21, 080009002A3E</code></p> <p>where the 12-digit hexadecimal number at the end of the message is the link level address of the LAN interface and is unique for each LAN interface.</p>	See 5.	See Flowchart H.

**(Flowchart A Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
5. Does the hostname of the diskless cnode's root server appear in the menu of bootable systems?	See Flowchart G.	See 6.
6. It is now time to check that the cluster configuration file, <code>/etc/clusterconf</code> , is correct. Go to Flowchart E. Return here after you have thoroughly examined <code>/etc/clusterconf</code> .	See 7.	See Flowchart E.
7. It is now time to verify that the remote boot daemon, <code>/etc/rbootd</code> , is running without error. Go to Flowchart B. Return here after you have thoroughly examined <code>/etc/rbootd</code> behavior.	Go to Flowchart F.	Go to Flowchart B.



**Flowchart B**

## **/etc/rbootd Not Running (Flowchart B)**

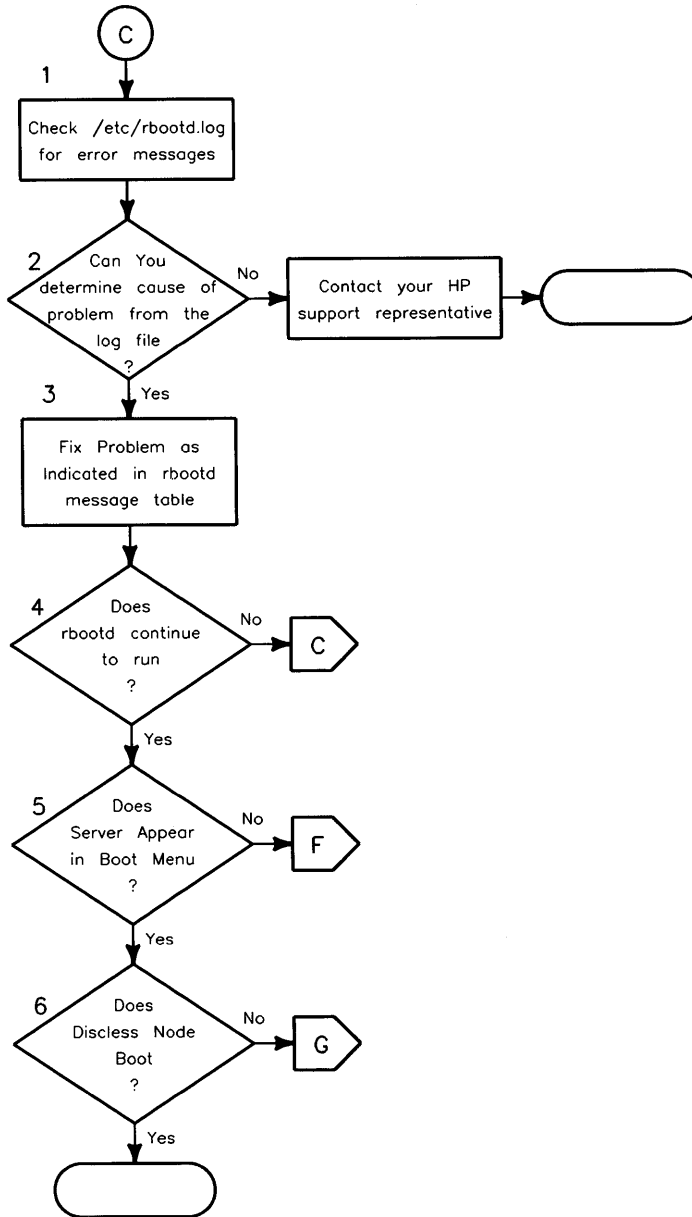
Use this flowchart to determine why the remote boot daemon, `/etc/rbootd`, is not running on the root server system.

<b>Question</b>	<b>Yes: Action</b>	
<p>1. Execute the following command: <code>-e   grep rbootd</code></p> <p>on the root servers to determine if the remote boot daemon program, <code>/etc/rbootd</code>, is running. If it is running, this command will return something like this:</p> <pre>118 ? 0:50 rbootd</pre>	See Flowchart D.	See 2.
<p>2. Try to manually start <code>/etc/rbootd</code>. Refer to the <code>rbootd(1M)</code> entry in the <i>HP-UX Reference</i>. Determine which LAN special device file should be used for <code>/etc/rbootd</code>. This should be the same LAN that the diskless cnode is on. All diskless cnodes of a root server must be on the same LAN. If no device file is specified, <code>/etc/rbootd</code> will use <code>/dev/ieee</code> as a default. Verify that the <code>/dev/ieee</code> device special file correctly describes the LAN interface used for cluster communication. Be sure that the device file has the correct major and minor numbers. It is also a good idea to execute <code>/etc/rbootd</code> with error logging, option <code>-13</code>, when you are troubleshooting. For instance, if you have determined that <code>/dev/ieee</code> is the special device file which should be used by <code>/etc/rbootd</code>, type the following command:</p> <pre>/etc/rbootd -13/dev/ieee</pre>	See 3.	See 3.

**(Flowchart B Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
3. Does <code>/etc/rbootd</code> continue to run? Repeat the <code>ps</code> command given in Step 1 above and determine if <code>/etc/rbootd</code> is still running.	See 4.	See Flowchart C.
4. Try to determine why <code>/etc/rbootd</code> was not running. <code>/etc/rbootd</code> is executed from <code>/etc/rc</code> . Be sure that <code>/etc/rbootd</code> is getting started when <code>/etc/rc</code> executes on the root server.	See 5.	See 5.
5. Go back to the diskless cnode. Do bootable systems for the root server's hostname now appear in the boot menu?	Make the appropriate menu selection to boot the diskless cnode. See 6.	See Flowchart D.
6. Did the diskless cnode boot successfully?	Troubleshooting is completed.	See Flowchart G.





**Flowchart C**

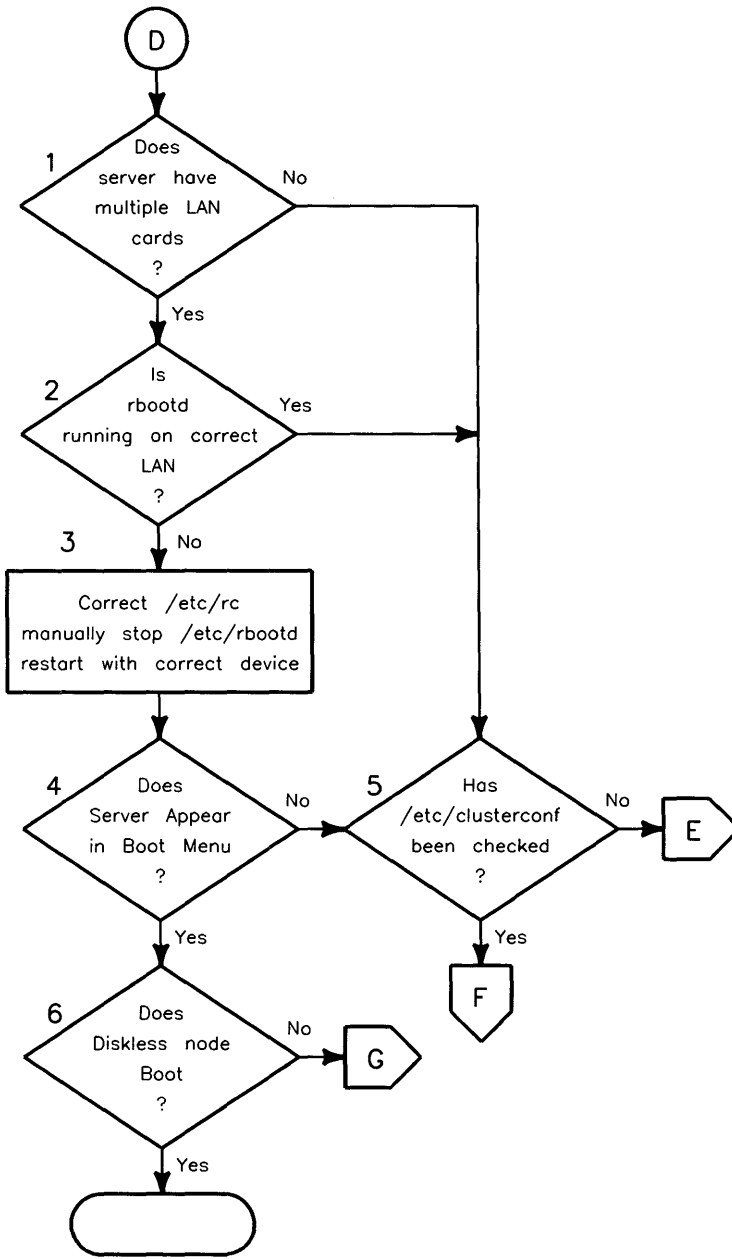
## **/etc/rbootd Not Running (Flowchart C)**

This flowchart is a continuation of /etc/rbootd troubleshooting. If you have not followed the steps in Flowchart B, go back to Flowchart B and start there.

Question	Yes: Action	No: Action
<p>1. /etc/rbootd logs information including error messages to a log file, /etc/rbootd.log. If you have not executed /etc/rbootd with logging, -l3 option, do this now. Be sure to run /etc/rbootd with the correct /dev LAN device special file. The correct format of the command is:</p> <pre data-bbox="154 646 494 670">/etc/rbootd -l3 /dev/ieee</pre> <p>Now check the /etc/rbootd.log file for messages. Try to determine the cause of your problem and how to solve the problem by consulting the rbootd error message table later in this troubleshooting section.</p>	See 2.	See 2.
<p>2. Were you able to determine why /etc/rbootd was unable to run? Could you fix the problem?</p>	See 3.	Contact your HP support representative for assistance.
<p>3. Be sure to make the fix permanent by modifying /etc/rc if necessary.</p>	See 4.	See 4.
<p>4. Use the ps command to verify that /etc/rbootd is continuing to run.</p>	See 5.	Repeat Flowchart C to see if there is a new error message in /etc/rbootd.log.
<p>5. Go back to the diskless cnode. Do bootable systems for the root server's hostname now appear in the menu?</p>	See 6.	See Flowchart F.

**(Flowchart C Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
6. Select the appropriate menu item and attempt to boot the diskless cnode. Does the cnode boot successfully?	Troubleshooting complete.	See Flowchart G.



Flowchart D

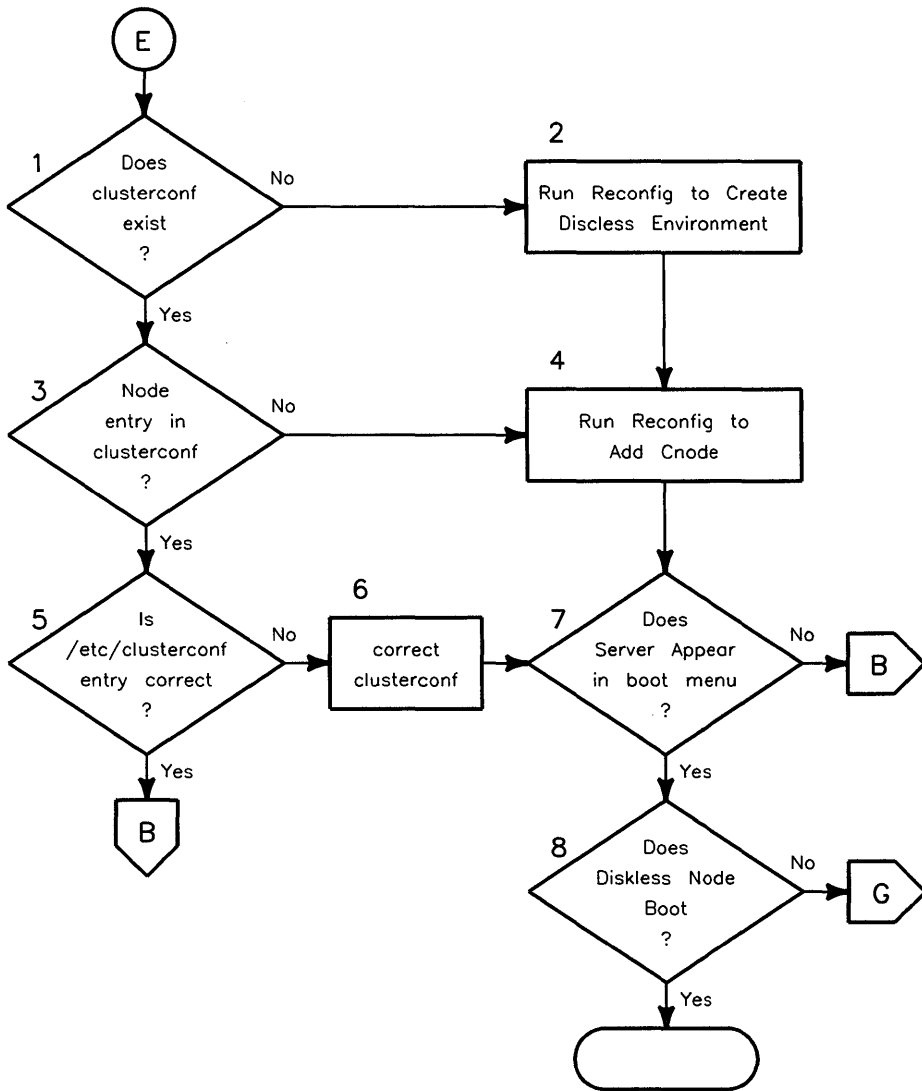
## Check for Multiple LAN Interfaces in Server (Flowchart D)

Use this flowchart if the root server has multiple LAN cards.

Question	Yes: Action	No: Action
<p>1. Does the root server have multiple LAN interfaces installed?</p> <p>Many HP 9000 Series 300 systems have a built-in LAN interface and may also have an additional HP 98643A card installed.</p>	See 2.	See 5.
<p>2. Verify that <code>/etc/rbootd</code> is running on the LAN device file which corresponds to the LAN the diskless cnodes for this server are on. All diskless cnodes must be on the same LAN. Is <code>/etc/rbootd</code> running on the correct device special file?</p> <p>You may have LAN device special files for more than one LAN interface. Check to see that <code>/etc/rbootd</code> is running on the LAN interface which is used for communication with this server's diskless nodes. Check the select code of this LAN interface card with the select code of the minor number of the LAN device special file used by <code>/etc/rbootd</code>. If no device file is specified, <code>/etc/rbootd</code> will use <code>/dev/ieee</code> by default.</p>	See 5.	See 3.
<p>3. Stop the currently running <code>/etc/rbootd</code> using the <code>kill</code> command. Restart <code>/etc/rbootd</code> with the correct device special file specified. You may either correct the <code>/dev/ieee</code> device special file or execute <code>/etc/rbootd</code> with another device name. The format of the <code>/etc/rbootd</code> command is:</p> <pre data-bbox="233 1349 532 1373">/etc/rbootd /dev/ieee</pre> <p>Make the change permanent by modifying <code>/etc/rc</code> if necessary.</p>	See 4.	See 4.

**(Flowchart D Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
4. Go back to the diskless cnode. Do bootable systems for the root server's hostname appear in the menu?	See 6.	See 5.
5. Have you already verified that /etc/clusterconf is correct by performing the steps in Flowchart E?	See Flowchart F.	See Flowchart E.
6. Make the appropriate menu selection and attempt to boot the diskless cnode. Does the cnode boot successfully?	Troubleshooting complete.	See Flowchart G.



Flowchart E

## **/etc/clusterconf Troubleshooting (Flowchart E)**

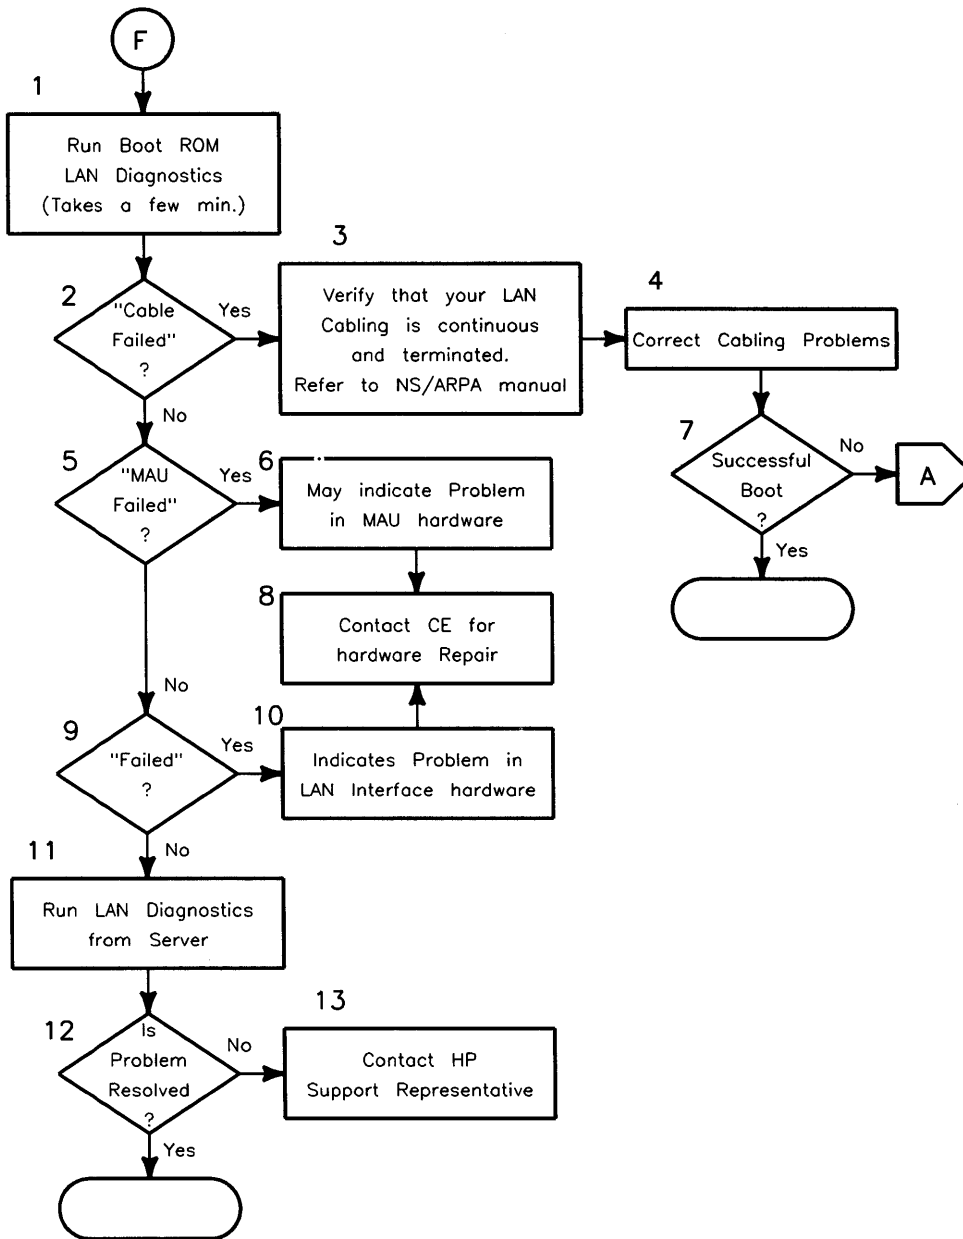
Use this flowchart to verify that `/etc/clusterconf` entries are correct.

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
<p>1. Verify that the <code>/etc/clusterconf</code> file exists and that the ownership and permissions of this file are correct. Permissions should be:</p> <pre style="margin-left: 40px;">-r--r--r--</pre> <p><code>/etc/clusterconf</code> should be owned by <code>root</code>, and the group should be <code>other</code>. Does <code>/etc/clusterconf</code> exist?</p>	See 3.	See 2.
<p>2. Create <code>/etc/clusterconf</code> by running the <code>reconfig</code> command if you have not already run <code>reconfig</code> to convert your system to a cluster root server system. Refer to “Configuring your HP-UX Cluster” section in the “Customizing the HP-UX System” chapter.</p> <p>If you have already completed this operation, but <code>/etc/clusterconf</code> has been removed from your system, try to recover the file from a system backup tape.</p>	See 3.	See 3.
<p>3. Look at the contents of <code>/etc/clusterconf</code> using the <code>more</code> command. Do you see an entry for the diskless cnode that cannot boot?</p>	See 5.	See 4.
<p>4. Create an entry in <code>/etc/clusterconf</code> for the diskless cnode by running the <code>reconfig</code> command to add a cnode. Refer to the “Configuring your HP-UX Cluster” section in the “Customizing the HP-UX System” chapter.</p>	See 7.	See 7.



**(Flowchart E Cont.)**

Question	Yes: Action	No: Action
<p>5. Refer to the <code>/etc/clusterconf(5)</code> entry in the <i>HP-UX Reference</i> and to the examples in the cluster configuration section of this manual. Is the <code>/etc/clusterconf</code> entry for this diskless cnode correct? You may also wish to run the <code>/etc/ccck</code> command which verifies the entries in the <code>/etc/clusterconf</code> file. Check the first field of the <code>/etc/clusterconf</code> entry to verify that it is the same as the 12-digit hexadecimal link address reported by the boot ROM identification messages on the diskless cnode's screen.</p>	See Flowchart B.	See 6.
<p>6. Correct the error in <code>/etc/clusterconf</code> entry. It may be that the link address was entered incorrectly. Correct this by running the <code>reconfig</code> command first deleting the cnode and then adding it back or you may use <code>vi</code> to edit <code>/etc/clusterconf</code>.</p>	See 7.	See 7.
<p>7. Go back to the diskless cnode. Do bootable systems for the root server's hostname now appear in the menu?</p>	See 8.	See Flowchart B.
<p>8. Select the appropriate menu item and attempt to boot the diskless cnode. Does the cnode boot successfully?</p>	Troubleshooting complete.	See Flowchart G.



**Flowchart F**

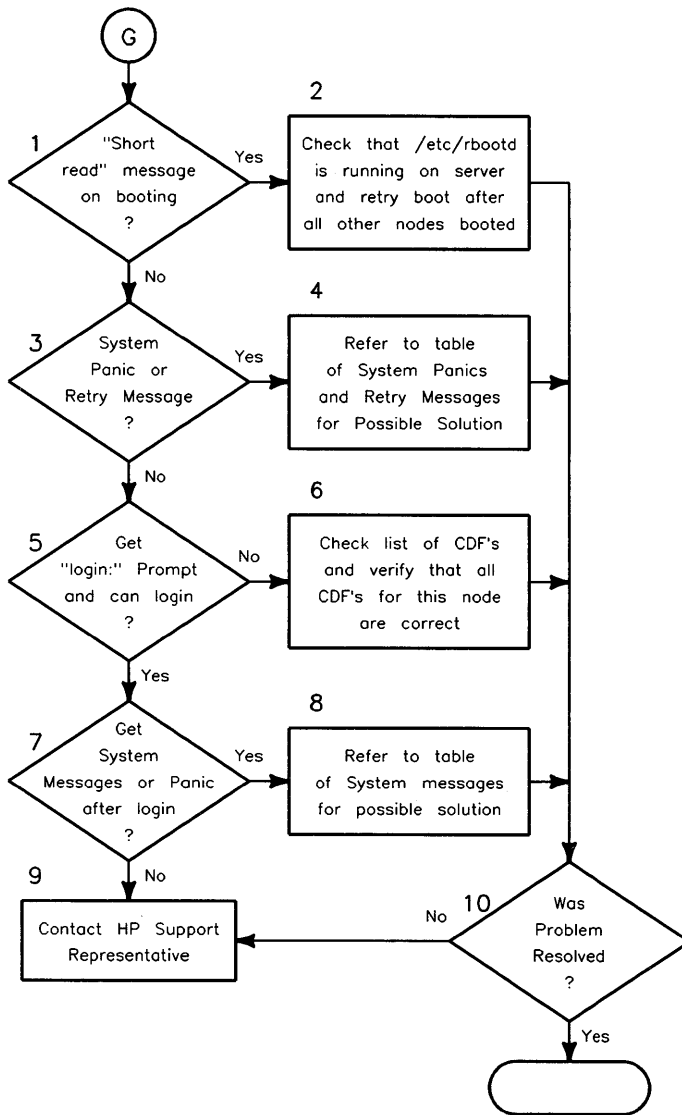
## Run Hardware Diagnostics (Flowchart F)

Use this flowchart to check for possible LAN hardware or cabling problems.

Question	Yes: Action	No: Action
1. Run the boot ROM LAN diagnostics on the diskless cnode. To do this be in attended mode and type ^C to get menu of extended tests. Type T for extended test. After extended memory test complete, check for a message following the 98643A device entry. If the extended test detected a hardware problem, you will see one of these messages: Cable failed, MAU failed, or Failed.	See 2.	See 2.
2. Did the extended test report Cable failed?	See 3.	See 5.
3. Cable failed indicates that there is a problem in the LAN cabling. Verify that the MAU is attached to the LAN cable and that the cable is terminated. Check for bad or broken cables or connectors. Refer to the <i>Installing and Maintaining NS-ARPA Services</i> manual for cable troubleshooting.	See 4.	See 4.
4. Correct any cabling problems that you have found and verify by trying the extended Boot ROM test again.	See 7.	See 7.
5. Did the extended test report MAU failed?	See 6.	See 9.

**(Flowchart F Cont.)**

Question	Yes: Action	No: Action
<p>6. MAU failed may indicate that there is a failure in the MAU LAN hardware. Since this may not be the case, try the MAU or LAN interface in another system that you know is working before you contact your HP support representative. Does the hardware work in another system?</p>	<p>See 11. If you were able to get the LAN hardware to work in another system, you may wish to go to step 11 and try to determine the cause of the problem by running additional diagnostics.</p>	<p>See 8.</p>
<p>7. Try to boot the system. Does the cnode boot successfully?</p>	<p>Troubleshooting complete.</p>	<p>Repeat troubleshooting from flowchart A.</p>
<p>8. Contact your HP hardware support representative for repair.</p>	<p>Troubleshooting complete.</p>	<p>Troubleshooting complete.</p>
<p>9. Did the extended test report Failed?</p>	<p>See 10.</p>	<p>See 11.</p>
<p>10. The Failed message indicates a problem in the LAN interface, the HP 98643A interface. Contact your HP hardware support representative for repair.</p>	<p>Troubleshooting complete.</p>	<p>Troubleshooting complete.</p>
<p>11. Refer to the LAN troubleshooting section in the <i>Installing and Maintaining NS-ARPA Services</i> manual for further software testing which may be run from the root server system.</p>	<p>See 12.</p>	<p>See 12.</p>
<p>12. Has the problem been resolved?</p>	<p>Troubleshooting complete.</p>	<p>Contact your HP support representative for assistance.</p>



**Flowchart G**

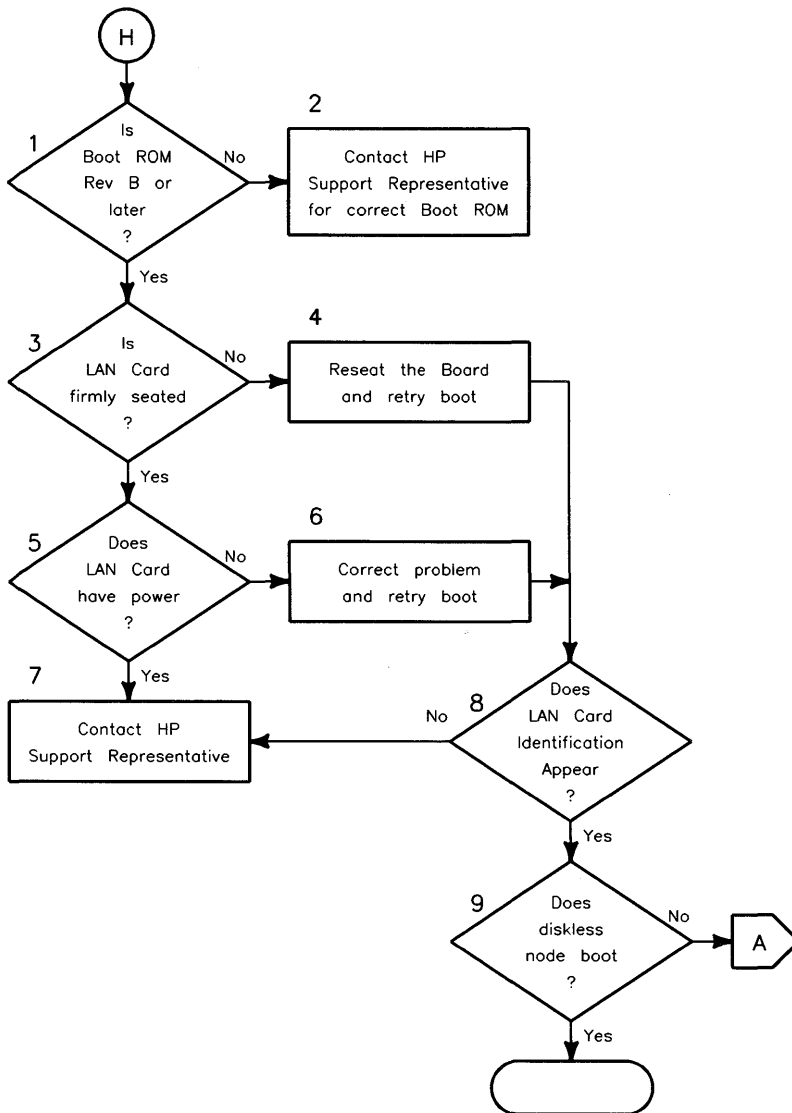
## Other Possible Problems (Flowchart G)

Use this flowchart if you are experiencing system panics or other messages when attempting to boot the diskless cnode or if the boot appears to complete, but you do not get a login prompt.

Question	Yes: Action	No: Action
1. Do you get the message <b>Short read</b> on the diskless cnode console when attempting to boot?	See 2.	See 3.
2. Check to see that <code>/etc/rbootd</code> is still running on the root server. This error could occur if <code>/etc/rbootd</code> terminated in the middle of the boot. Refer to Flowchart B for details.  This error might also occur if you are trying to boot more than the number of supported diskless cnodes at the same time. If this is the case, try to boot the failed cnode after the other cnodes have finished booting.	See 10.	See 10.
3. Does the diskless cnode get a panic or retry message when attempting to boot?	See 4.	See 5.
4. Refer to the table of "Possible panic, retry and other system messages" later in this appendix for an explanation of the cause and possible fix for your problem.	See 10.	See 10.
5. Do you get the <code>login:</code> prompt and can you successfully login to the diskless cnode?	See 7.	See 6.

**(Flowchart G Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
6. Check the list of CDF's in the "Cluster Concepts" section of the "System Management Concepts" chapter of this manual. Verify that all the CDF's were created correctly for the cnode which does not boot or give a <b>login:</b> prompt. In particular, be sure to check that <code>/etc/inittab</code> , <code>/dev</code> and <code>/etc/wtmp</code> CDF's were created correctly for this cnode. Correct any problems you find.	See 10.	See 10.
7. Do you get a system panic or a system message some time after the diskless cnode boots and you login to the cnode? The system message may appear on the root server as well as a diskless node.	See 8.	See 9.
8. Refer to the table of "Possible panic, retry and other system messages" later in this appendix for the cause and a possible fix for the problem.	See 10.	See 10.
9. You have been unable to determine the cause of the problems you are having with your HP-UX cluster from this troubleshooting procedure. For further assistance, contact your HP support representative.	Troubleshooting complete.	Troubleshooting complete.
10. Were you able to determine the cause of the problem?	Troubleshooting complete.	See 9.



**Flowchart H**



## Some Additional Hardware Checks (Flowchart H)

Question	Yes: Action	No: Action
1. Is the Boot ROM in your diskless computer Revision B or later? Refer to the "System Startup and Shutdown" chapter of this manual for instructions on determining what revision boot ROM your system has.	See 3.	See 2.
2. A revision B or later version of the boot ROM is required to be able to boot a diskless node. Contact your HP support representative for information on obtaining the correct revision boot ROM.	Troubleshooting complete.	Troubleshooting complete.
3. If there is no LAN interface, HP 98643A, diagnostic displayed on the left side of the screen, be sure that the LAN card is installed and that it is properly seated in the card cage.	See 5.	See 4.
4. Reseat the board. If necessary, try the board in a different slot in the card cage.	See 8.	See 8.
5. Does the LAN card have power? If the LAN card is in an expander unit, is the expander powered on?	See 7.	See 6.
6. Correct the problem so that power is supplied to the LAN card.	See 8.	See 8.
7. If all hardware appears to be properly installed and the correct boot ROM is installed in your system, contact your HP support representative for assistance. You may first wish to try the suspect LAN interface in a different system to be sure that it is the interface and not some other part of the system that is at fault.	Troubleshooting complete.	Troubleshooting complete.

**(Flowchart H Cont.)**

<b>Question</b>	<b>Yes: Action</b>	<b>No: Action</b>
8. Now retry the attended boot and check for the correct LAN diagnostic as described in step 4 of Flowchart A. Does the correct message appear?	See 9.	See 7.
9. Attempt to boot the diskless cnode. Does the cnode boot successfully?	Troubleshooting complete.	See Flowchart A.

## Possible Boot Daemon Error Messages

Use this table to resolve the cause of `/etc/rbootd` error messages. These messages are logged in the file `/etc/rbootd.log` if `/etc/rbootd` is run with logging enabled.

Error Message	Cause/Fix Information
<p>Could not expand ???????????????? for client ??????????????.</p>	<p><code>rbootd</code> could not resolve a context dependent file (CDF). Check that the <code>/hp-ux</code> CDF has been created correctly and that a version exists for the cnode.</p>
<p>Could not open <code>/dev/null</code>.</p>	<p>Check to be sure that <code>/dev/null</code> exists. Verify that permissions on <code>/dev/null</code> and <code>/dev</code> are correct.</p>
<p>Could not open boot configuration file/<code>/etc/boottab</code></p>	<p><code>rbootd</code> could not open <code>/etc/boottab</code>. Verify that <code>/etc/boottab</code> exists and that the file has correct permissions.</p>
<p>Could not open lan device ??????????</p>	<p><code>rbootd</code> could not open the lan device specified by the device special file ??????????. Verify that the device special file given has the correct major and minor numbers. If the device file is not correct, remake the device file or execute <code>rbootd</code> with another device file specified in the command line. <code>/etc/rbootd</code> uses <code>/dev/ieee</code> as default. If another LAN device special file should be used, change the <code>/etc/rbootd</code> command to specify the device: <code>/etc/rbootd /dev/thicklan</code></p> <p>It is not necessary to reboot. Just type the <code>/etc/rbootd</code> command. You should be sure to modify <code>/etc/rc</code> so <code>/etc/rbootd</code> will be executed correctly following any future reboots.</p>
<p>Could not stat boot configuration file <code>/etc/boottab</code></p>	<p><code>rbootd</code> could not access the <code>/etc/boottab</code> file. Verify that <code>/etc/boottab</code> exists.</p>

## (/etc/rbootd Error Messages Cont.)

Error Message	Cause/Fix Information
Could not stat cluster configuration file /etc/clusterconf.	rbootd could not open /etc/clusterconf. Verify that /etc/clusterconf exists. If /etc/clusterconf does not exist, refer to "Configuring your HP-UX Cluster" in the chapter "Customizing the HP-UX System" of this manual for instructions on setting up a cluster.
Hostname not set.	The hostname of the server system is not set. Check /etc/rc to be sure that hostname is getting set correctly. After correcting the problem, either reboot the server or set the hostname with the hostname command.
Illegal link address 080009?????? in boot configuration file. Entry ignored.	rbootd found an invalid entry in /etc/clusterconf. The first field in the entry is incorrect. This field is the ethernet address and is of the format 080009?????? where ?????? is a unique 6-digit hexadecimal number. Check /etc/clusterconf and correct the error.
Out of session space.	rbootd can only process boot requests from 50 systems at any time. This message is logged when more than 50 systems are requesting to boot from the server. Wait until some systems have booted and retry the boot.
Received request for closed session.	rbootd has received a request from a boot session which rbootd had recorded as terminated. Retry the boot.

### (/etc/rbootd Error Messages Cont.)

Error Message	Cause/Fix Information
Rejected boot request from 080009???????	rbootd has rejected the boot request from a system with ethernet address 080009??????. Check /etc/clusterconf to verify that an entry exists for this cnode and that the entry is correct. Check that /etc/boottab has not been corrupted. Check /hp-ux CDF to verify that a version exists for this cnode.
Requested boot file (SYSHP,SYSBACKUP, SYSDEBUG /hp-ux) is not an ordinary file.	The boot file requested by the diskless cnode is not an ordinary file. Verify that the permissions of /usr/boot and /hp-ux are correct. Check the version of the /hp-ux CDF for this cnode to verify that the permissions are correct.

## Possible Panic, Retry, and Other System Messages on Diskless Cnodes

Refer to this table if you are experiencing diskless cnode panics, retry messages on the diskless cnode or other system messages on either the root server or a diskless cnode. If a panic is occurring, the cause of the problem should first be resolved. Then the cnode will have to be rebooted by cycling power. If a retry message is occurring, the diskless cnode will continue to retry the boot until the cause of the problem has been fixed. Once the cause of the problem has been fixed, the boot should continue.

The messages in the following table are listed in alphabetical order.

Message	Type	Cause/Fix Information
cannot exec /etc/read_cct	Retry	During boot of a diskless cnode, /etc/read_cct is executed to read the /etc/clusterconf entry. Verify that /etc/read_cct exists and that the permissions on the file are correct. The permissions should be: <code>-r-wr--r-- root other</code>
Client has NFS and server does not. Reconfigure/hp-ux.	Panic	In a cluster, NFS must be configured for all cnodes or NFS must not be configured for any cnode. In this case, the diskless cnode kernel contains the NFS code and the server kernel does not. Either make a new diskless cnode kernel removing NFS, or make a new server kernel adding NFS.

## (System Messages Cont.)

Message	Type	Cause/Fix Information
Client kernel's revision incompatible with server's( <code>uname -r</code> )	Panic	The revision code of the diskless cnode kernel does not match that of the root server. Check the <code>/hp-ux</code> CDF files for obsolete kernel files. Use the <code>what</code> command to get the <code>what</code> string for the diskless cnode versions of <code>/hp-ux</code> and compare this with the result of <code>uname -r</code> on the root server. The kernel revision numbers should match. If not, make a new kernel for the server or diskless cnode, whichever is older.
CSP shortage on Server	Retry	There are insufficient cluster server processes on the root server for the diskless cnode to boot. Start additional cluster server processes on the server.
Diskless: LAN Failure, Invalid Card State	Panic or System Message	See discussion of possible LAN problems following this table.
Diskless: LAN Failure, Unknown Cause	Panic or System Message	See discussion of possible LAN problems following this table.
Diskless: LAN Hardware Failure	Panic or System Message	See discussion of possible LAN problems following this table.
Diskless: LAN Interface Card Failure	Panic or System Message	See discussion of possible LAN problems following this table.
Diskless: LAN Link Failure	Panic or System Message	See discussion of possible LAN problems following this table.

## (System Messages Cont.)

Message	Type	Cause/Fix Information
<code>/etc/clusterconf: cnode idout of bounds</code>	Retry	The entry for this cnode in <code>/etc/clusterconf</code> , is missing or incorrect. Check the file <code>/etc/clusterconf</code> . If the entry does not exist, see the section in this manual on cluster configuration. If <code>/etc/clusterconf</code> exists, verify that the entry for the failed cnode is correct. The maximum cnode id is 255. Verify that the cnode is less than 255 and greater than or equal to 1. The command <code>ccck</code> may be used to verify the correctness of the <code>/etc/clusterconf</code> file. Refer to the <i>ccck(1M)</i> entry in the <i>HP-UX Reference</i> .
<code>/etc/clusterconf: duplicate cnode id</code>	Retry	When this diskless cnode attempted to boot, a duplicate cnode id was detected. Check <code>/etc/clusterconf</code> on the server and fix the duplicate cnode id problem. The command <code>ccck</code> may be used to verify the correctness of the <code>/etc/clusterconf</code> file. Refer to the <i>ccck(1M)</i> entry in the <i>HP-UX Reference</i> .
<code>/etc/read_cct aborted</code>	Retry	The command <code>/etc/read_cct</code> is a program that reads the <code>/etc/clusterconf</code> file on the root server while a diskless cnode is joining the cluster. This message indicates that <code>/etc/read_cct</code> has terminated prematurely. Check that the <code>/etc/read_cct</code> is the correct version for the version of the operating system that you are running. If you are unable to resolve the problem, contact your HP support representative.



**(System Messages Cont.)**

Message	Type	Cause/Fix Information
Failed kernel selftest: Cannot allocate file system buffer.	System Message	<p>NOTE: IGNORE THIS MESSAGE IF IT IS ACCOMPANIED BY THE MESSAGE: Cannot invoke cluster server process</p> <p>The cnode cannot allocate file system buffers. Refer to the "Operating System Parameters" appendix in this manual and make a new kernel that increases the value of nbuf.</p>
Failed kernel selftest: Cannot allocate kernel message buffer.	System Message	<p>The cnode has run out of message buffers. Refer to the "Operating System Parameters" appendix in this manual. Make a new kernel increasing the value of dskless_mbufs and dskless_cbufs.</p>
Failed kernel selftest: Cannot allocate fs_bufs.	System Message	<p>The cnode cannot allocate file system buffers under interrupt. Refer to the "Operating System Parameter" appendix in this manual to determine how to increase the dskless_fsbufs parameter and make a new kernel for this cnode.</p>
Failed kernel selftest: Cannot invoke cluster server process	System Message	<p>Indicates that the cluster server processes are being tied up. If this message repeats often, try starting additional general cluster server processes. Refer to the <i>csp(1M)</i> entry in the <i>HP-UX Reference</i>. Be sure to make any change permanent by updating the <code>/etc/clusterconf</code> file. NOTE: You may need to reconfigure the diskless cnode kernel before starting general cluster server processes on a diskless cnode. The kernel parameter you need to modify is <code>ngcsp</code>. This kernel parameter should be 0 for diskless cnodes and at least twice the number of cluster server processes you plan to run on the diskless cnode.</p>

**(System Messages Cont.)**

Message	Type	Cause/Fix Information
FATAL ERROR: DISKLESS LAN FAILURE: Card State = ?	Panic or System Message	See discussion of possible LAN problems following this table.
not in /etc/clusterconf	Retry	The /etc/read_cct program could not find the diskless cnode in the /etc/clusterconf file or /etc/clusterconf is missing. See Flowchart E for troubleshooting /etc/clusterconf problems.
Server has NFS and client does not. Reconfigure/hp-ux.	Panic	In a cluster, NFS must be configured for all cnodes or NFS must not be configured for any cnode. In this case, the server kernel contains the NFS code and the diskless cnode kernel does not. Either make a new diskless cnode kernel adding NFS, or make a new server kernel removing NFS.
Server is not a rootserver	Panic	/etc/rbootd (the remote boot daemon) is running on a system which is not a root server system. Determine why rbootd was running. /etc/rbootd is executed from /etc/rc.
Server not responding (clustered?)	Retry	rbootd (the remote boot daemon) is running on a system, but the system has not clustered. Verify that the system should be a root server. If it should not be a root server, determine why /etc/rbootd was running and correct. /etc/rbootd is executed from /etc/rc. If the system should be a root server, determine why the system is not clustered. /etc/cluster is executed from /etc/rc.

**(System Messages Cont.)**

Message	Type	Cause/Fix Information
Suspected AUI cable disconnected from MAU or grounded backbone cable	System Message	<p>This message may also appear a few times before successful boot in a VERY busy cluster, but the boot should continue in this case.</p> <p>See discussion of possible LAN problems following this table.</p>
Suspected backbone cable not properly terminated	System Message	<p>See discussion of possible LAN problems following this table.</p>
Suspected backbone cable not properly terminated or MAU disconnect.	System Message	<p>See discussion of possible LAN problems following this table.</p>
swapconf - no swap area associated with entry 0	Panic	<p>This message will be preceded by one of the following messages:</p> <p>1) <b>ignored - can't open device</b> The diskless node is expecting to run with a local swap disk. This message is caused by the swap disk not being powered on or the driver for the disk not being configured in the kernel or the swap device not being configured in the kernel. Refer to "Local and Remote Swapping in an HP-UX Cluster" in the "Kernel Customization" chapter of this manual for details on configuring a diskless cnode kernel for a local swap device.</p> <p>2) <b>ignored - can't open file system</b> The diskless cnode is expecting to run with a local swap disk. This message may appear if the kernel is expecting to swap on a disk that has a file system and there is no file system on the disk.</p>

## System Messages Cont.)

Message	Type	Cause/Fix Information
		Refer to “Local and Remote Swapping in an HP-UX Cluster” in the “Kernel Customization” chapter of this manual for details on configuring a diskless cnode kernel for a local swap device. 3) ignored - no room The diskless cnode is expecting to run with a local swap disk. This message appears if the disk being used for local swap was initialized with no swap space on the disk. Refer to “Local and Remote Swapping in an HP-UX Cluster” in the “Kernel Customization” chapter of this manual for details on initializing disks for swap and configuring a kernel for a swap device.

### LAN Cable Break Detection Behavior

The HP-UX cluster protocol, in conjunction with the diskless recovery code, is frequently capable of surviving a broken or unterminated LAN cable<sup>2</sup>. However, there are LAN cable topological configurations that must be considered prior to configuring a cluster. The cluster LAN break detection and recovery code will not detect a broken or unterminated LAN if the diskless cnodes and their respective root server cnode are on opposite sides of either a LAN bridge box, LAN repeater, or any other device that acts as a terminator (Fig 1). In addition, if the MAU or AUI cable is disconnected from a diskless cnode, then the root server, after certain check alive periods, will probably declare the cnode dead. This situation is only detectable on the diskless cnode in question; the rest of the LAN is still functioning correctly. The converse situation is where the root server’s MAU or AUI cable is disconnected. This will most likely result in all the diskless cnodes losing contact with the root server.

<sup>2</sup> It is recommended that the entire cluster be shut down prior to performing any LAN cable maintenance.

If the LAN cable is broken or unterminated the following messages may be received. These messages are considered recoverable temporary failures:

`Suspected backbone cable not properly terminated.`

`Suspected backbone cable not properly terminated or MAU disconnect.`

`Suspected AUI cable disconnected from MAU or grounded backbone cable.`

The following messages may be received and are non-recoverable failures on the diskless cnode:

■ `Diskless: LAN Failure, Unknown Cause`

■ `FATAL ERROR: DISKLESS LAN FAILURE: Card State = X`

(where X is a number that is interpreted for you as one of the following panics).

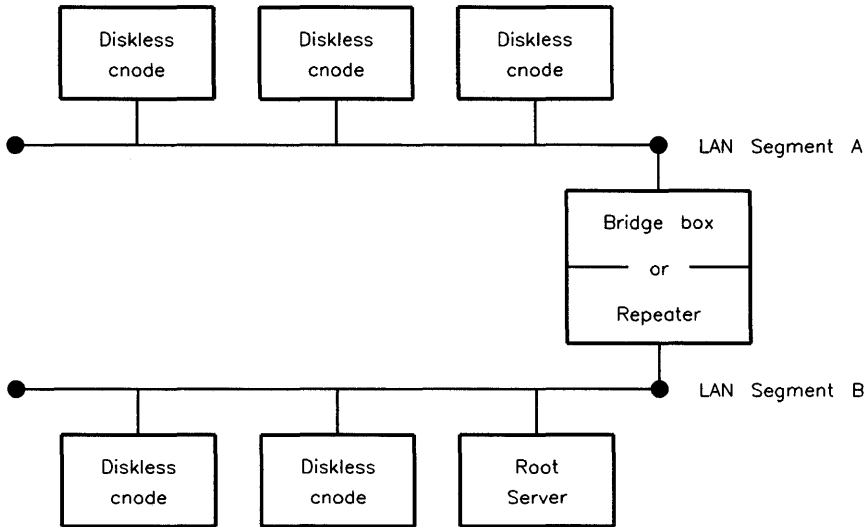
□ `Diskless: LAN Interface Card Failure`

□ `Diskless: LAN Link Failure`

□ `Diskless: LAN Hardware Failure`

□ `Diskless: Lan Failure, Invalid Card State`

If these errors occur on the root server, the LAN device driver is left in a DOWN state and all cnodes will lose contact with the root server cnode.



**Figure B-1. Bridgebox Setup**

If segment B were broken or unterminated, the diskless cnodes (W) on segment A would lose contact with their root server (S). The diskless cnodes and root server on segment B would recognize the problem and wait until the LAN is repaired. There may be a slight delay after the LAN is repaired for the recovery code to declare the LAN as being UP and for transmissions/receptions to continue.

## Unable to Login or Rlogin to a Diskless Cnode

If you can boot the diskless cnode, but cannot log into the diskless cnode, you probably have one or more bad system files that are unique to that cnode. If you cannot boot the diskless cnode, refer to the previous section.

Possible problem files are:

```
/etc/inittab+ /dev+    /etc/ioctl.syscon+  
/hp-ux+          /etc/checklist+ (only  
                  if you have local swap)
```

One probable file is the `/etc/inittab` file. To edit the cnode's version of the file, change to the CDF directory, then access the file with the same name as the cnode that cannot boot (refer to the "System Management Concepts" chapter, section "Cluster Concepts" for a full description of CDFs and how to access the files within CDFs). For example, to edit the `/etc/inittab` file without logging into the cnode called donald, type:

```
cd /etc/inittab+  
vi donald
```

If the problem is with the kernel, `/hp-ux`, then a new kernel must be created. *You must use the `configcommand` to create a kernel for another cnode. Refer to the chapter "Kernel Customization" for more information on the `configcommand`. Do not use the `reconfigprogram` to create a kernel for a cnode unless you can execute it from the cnode itself.* When you use the `config` program to create a kernel, build the kernel in a temporary location (*NOT* `/hp-ux`), then type the following to set up the `/hp-ux` and kernel description files for the cnode:

```
mv hp-ux /hp-ux+/<cnodename>  
mv dfile /etc/conf/dfile+/<cnodename>
```





# Contents

---

## C. Partitions and Filesets

How Partitions and Filesets Work . . . . .	C-1
Dependent Filesets . . . . .	C-2
Mandatory vs. Non-mandatory Partitions . . . . .	C-2
Which Filesets Do I Want? . . . . .	C-3
Filesets and Their Functions . . . . .	C-4
Filesets, Dependencies, and Sizes . . . . .	C-9
Recovering Lost Files . . . . .	C-15
Recovering Files Required For Update . . . . .	C-16
If you installed from a cartridge tape: . . . . .	C-16
If you installed from micro disks: . . . . .	C-16
How to List Files . . . . .	C-17
List of AXE System Files . . . . .	C-18
Partition SYS_CORE . . . . .	C-18
Fileset ACAT . . . . .	C-18
Fileset ACORE . . . . .	C-18
Fileset APAM . . . . .	C-23
Fileset CORE . . . . .	C-23
Fileset KERNEL . . . . .	C-23
Fileset SSIT . . . . .	C-24
Fileset TOOL . . . . .	C-24
Partition TEXT . . . . .	C-24
Fileset AMANUAL . . . . .	C-24
Fileset ATEXT . . . . .	C-26
Partition SYS_TOOLS . . . . .	C-26
Fileset ABCMD . . . . .	C-26
Fileset ACMD . . . . .	C-27
Partition STARBASE . . . . .	C-28
Fileset AFA_FM . . . . .	C-28

Fileset	ASTARBAS . . . . .	C-29
Partition	X11WINDOWS . . . . .	C-29
Fileset	AGRM . . . . .	C-29
Fileset	AX11MAN . . . . .	C-29
Fileset	AX11RENV . . . . .	C-30
Fileset	AX11SERV . . . . .	C-31
Partition	WINDOW . . . . .	C-33
Fileset	AWINDOW . . . . .	C-33
Partition	PROG_LANGS . . . . .	C-34
Fileset	AC . . . . .	C-34
Fileset	AFPA . . . . .	C-34
Fileset	APROG . . . . .	C-34
Partition	MISC_UTILS . . . . .	C-34
Fileset	AAUDIT . . . . .	C-34
Fileset	ACONFIG . . . . .	C-34
Fileset	ATRMINFO . . . . .	C-36
Fileset	LAN . . . . .	C-44
Partition	NLS . . . . .	C-45
Fileset	NAMERICA . . . . .	C-45
Fileset	NCFRENCH . . . . .	C-45
Fileset	NCHINESS . . . . .	C-45
Fileset	NCHINEST . . . . .	C-45
Fileset	NDANISH . . . . .	C-45
Fileset	NDUTCH . . . . .	C-45
Fileset	NENGLISH . . . . .	C-45
Fileset	NFINNISH . . . . .	C-45
Fileset	NFRENCH . . . . .	C-46
Fileset	NGERMAN . . . . .	C-46
Fileset	NGREEK . . . . .	C-46
Fileset	NICELAND . . . . .	C-46
Fileset	NITALIAN . . . . .	C-46
Fileset	NJAPANES . . . . .	C-46
Fileset	NKOREAN . . . . .	C-46
Fileset	NNORWEGI . . . . .	C-46
Fileset	NPORTUGU . . . . .	C-46
Fileset	NSPANISH . . . . .	C-47
Fileset	NSWEDISH . . . . .	C-47
List of PE System Files	. . . . .	C-48

Partition	SYS_CORE	C-48
Fileset	FTN_CORE	C-48
Fileset	PAS_CORE	C-48
Fileset	PCORE	C-48
Fileset	PDSKLESS	C-49
Partition	TEXT	C-49
Fileset	HPUX_MAN	C-49
Fileset	PTEXT	C-78
Partition	SYS_TOOLS	C-79
Fileset	PCEUTIL	C-79
Fileset	PDRIVERS	C-82
Fileset	PFILTER	C-82
Fileset	PRLP	C-83
Fileset	PSYSCOM	C-83
Partition	STARBASE	C-84
Fileset	DSTARBAS	C-84
Fileset	MSTARBAS	C-84
Fileset	PFA_FM	C-88
Fileset	PSTARBAS	C-88
Partition	X11WINDOWS	C-89
Fileset	PX11MAN	C-89
Fileset	PX11PENV	C-92
Partition	WINDOW	C-95
Fileset	DWINDOW	C-95
Fileset	MWINDOW	C-96
Fileset	PWINDOW	C-98
Partition	PROG_LANGS	C-98
Fileset	PC	C-98
Fileset	PFORTRAN	C-102
Fileset	PPROG	C-102
Partition	MISC_UTILS	C-102
Fileset	PACCT	C-102
Fileset	PAPPLIC	C-103
Fileset	PPERIPH	C-104
Fileset	PSCCS	C-104
Partition	NLS	C-104
Fileset	NCORE	C-104
Fileset	NNCOMPUT	C-104

Fileset PORT\_NLS . . . . . C-106

## Partitions and Filesets

---

HP-UX (or AXE) is divided into pieces called *partitions*. Each partition is further divided into pieces called *filesets*. File sets are made up of HP-UX files (a set of files). You can load entire partitions (containing one or more filesets), or you can load individual filesets. File sets are the smallest unit in installation.

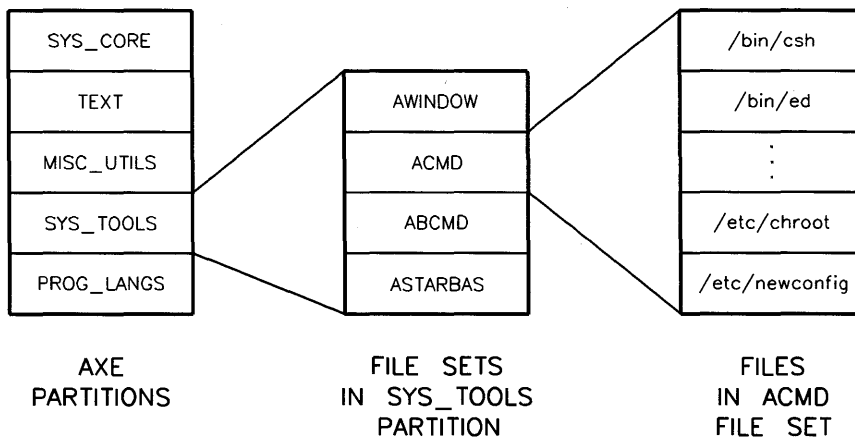
By separating HP-UX into filesets, you can install only those filesets you will use, thus saving time and disk space. For example, if you do not plan to use a part of HP-UX (such as files to support non-HP terminals) you have the option of not loading the fileset containing it (ATRMINFO).

Except for guarding against the possibility of trying to install partitions or filesets which will not fit on your hard disk, not much planning is required. In the installation procedure you have a great deal of flexibility in what you want on (or off) the system. If you do not load all of the partitions and filesets that you need during installation, they can be added later via the `/etc/update` process. Likewise, filesets can also be removed via the `sysrm` command.

---

### How Partitions and Filesets Work

You can think of partitions, filesets, and files as “building blocks” of the HP-UX system. The smallest unit of information in this scheme is a *file*. Groups of related files are grouped into *filesets*. File sets, in turn, are grouped into *partitions*. Finally, the partitions—taken as a whole—comprise the HP-UX system. Figure C-1 illustrates the relationship between the HP-UX system, partitions, filesets, and files:



**Figure C-1. HP-UX Partitions, Filesets and Files**

## Dependent Filesets

Some filesets require commands located in another fileset. The fileset is said to be *dependent* on the other fileset. If such a dependency exists, the fileset will be loaded along with (either before or after) the fileset in question.

When trying to load a fileset that has a dependency, the installation or update process informs you of the dependency and loads the dependency. If necessary, the installation or update process prompts you to insert a different disk from which to load; it specifies which disk or tape to use. If more than one dependency exists, the computer may prompt for yet another disk to use for loading.

File sets are never split across disks although a given fileset may have dependencies on other disks.

## Mandatory vs. Non-mandatory Partitions

Only one partition *must* be installed for your system to function: **SYS\_CORE**. The installation procedure automatically loads this partition when you install AXE (the core, minimum system). This partition contains the **ACORE** filesets which contain essential facilities and commands used by other filesets. All other partitions contain optional filesets (filesets that are not required for HP-UX to run).

All the other partitions are considered non-mandatory; they are not required to run HP-UX. Refer to the section on “Which Filesets Do I Want” for a discussion on choosing filesets.

---

## Which Filesets Do I Want?

To help you find the disk containing the desired fileset, the HP-UX system is partitioned into color-coded disks (if you received HP-UX in a set of micro disks). Each color represents one partition. Table C-1 shows the colors associated with each partition. If you received AXE on cartridge tape, all filesets and partitions are contained on one tape.

**Table C-1. Table Flexible Disk Label Colors**

Partition	Label Color
SYS_CORE	No colored label. Referred to as “grey” disks.
SYS_TOOLS	red
PROG_LANGS	blue
MISC_UTILS	green
TEXT	orange
GRAPHICS	purple

In addition to the mandatory `SYS_CORE` partition, you should install the `SYS_TOOLS` partition. `SYS_TOOLS` is not a mandatory partition, but it contains commands and tools necessary for file management, file manipulation, process control, peripheral support, and HP Windows/9000. The `SYS_TOOLS` partition is on the flexible disks with the red labels.

## Filesets and Their Functions

Besides these two partitions, the loading of all other partitions and filesets depends on your intended use of the system. You need to determine which filesets you want on the system. This factor must be weighed against your other needs for available storage space on your hard disk.

To assist you in selecting the partitions and filesets that you will need, Table C-2 and Table C-3 list possible tasks you will do on your system, along with the fileset required to execute the command.

---

<b>Note</b>	Before loading filesets, be sure to check that you have enough space on your file system. Refer to the next section, “Filesets, Dependencies, and Sizes” for more information.
-------------	--

---



**Table C-2. Application Execution Environment Functions**

Task	Fileset(s)	Notes
Use the Bourne shell. Back up your system.	CORE	Mandatory. This fileset is automatically loaded when you install AXE.  To perform a backup, refer to the section “Backing up and Restoring the File System” in this manual.
Use the PAM shell	APAM	This is a mandatory fileset and is automatically loaded when you install AXE.
Install optional software products. Update your system to another version.	TOOL	Mandatory. This fileset is automatically loaded when you install your system. To perform the installation or update, use the procedures in “Installing Optional Software and Updating Your HP-UX Kernel” in this manual.
Edit files, <code>pr</code> command, C shell	ACMD	Refer to <i>A Beginner’s Guide to Text Editing</i> for information on how to edit files, and <i>A Beginner’s Guide to Shells</i> for information on the C shell.

**Table C-2. Application Execution Environment Functions  
Continued**

Task	Fileset(s)	Notes
Reconfigure the kernel drivers, system parameters, or swap space (with either <code>reconfig</code> or <code>config</code> ).	ACONFIG	Automatically loads dependent filesets. The procedure to reconfigure your kernel drivers and swap space is in the section "Using the <code>config</code> Command" of Chapter 6 in this manual
Run HP Windows/9000	AWINDOW	Automatically loads dependent filesets.
Run applications that use graphics.		Automatically loads dependent filesets.
Use the floating point accelerator capabilities.	AFPA	Loads code required for the floating point accelerator onto your system. You must still reconfigure your kernel and load some microcode as explained in the installation note for the floating point accelerator.
BIF and LIF utilities, for example to list files on your installation or update media.	ABCMD	
Use the on-line documentation for AXE commands.	AMANUAL	Automatically loads dependent filesets.
Use non-HP terminals with your system.	ATRMINFO	

**Table C-3. Programming Environment Functions**

Task	Fileset(s)	Notes
Use Programming Environment (PE).	PCORE	This is a mandatory fileset, but is not automatically loaded. It automatically loads other dependent filesets.
Miscellaneous utilities such as <code>chsh</code> , <code>cal</code> , <code>lock</code> , <code>primes</code> , <code>shl</code> , <code>tic/untic</code> .	PAPPLIC	
C programming	PC	Automatically loads dependent filesets.
Miscellaneous utilities such as <code>adjust</code> , <code>unique</code> , <code>head</code> , <code>strings</code> , <code>pack/unpack</code> .	PFILTER	Automatically loads dependent filesets.
Use <code>mt</code> or <code>lsdev</code> .	PPERIPH	
Use programming development tools such as <code>adb</code> , <code>cdb</code> , <code>nm</code> , <code>yacc</code> , <code>lex</code> , <code>csplit</code> . Also use the <code>curses</code> library, device I/O library, and <code>malloc</code> library.	PPROG	Automatically loads dependent filesets. Tutorials for most of these tools are in the <i>Concepts and Tutorials</i> series of manuals.
Set up and use <code>uucp</code> , VT, <code>mailx</code> , <code>kermit</code> , and X.25 data communication features.	PSYSCOM	Automatically loads dependent filesets. Tutorials for most of these tools are in the <i>Concepts and Tutorials</i> series of manuals.
Text processing: <code>nroff</code> , <code>tbl</code> , <code>neqn</code> , <code>mm</code> , <code>spell</code> , <code>ms</code> , <code>fixman</code> , <code>mkwhatis</code> , <code>diffmk</code> , <code>catman</code> , <code>hyphen</code> .	PTEXT	Automatically loads dependent filesets.

**Table C-3. Programming Environment Functions  
Continued**

Task	Fileset(s)	Notes
Use on-line manual pages.	HPUX_MAN	Automatically loads dependent filesets.
Program with Starbase graphics.	DSTARBAS MSTARBAS PSTARBAS	Starbase demos. Starbase on-line manual pages. Starbase libraries for program development. Automatically loads dependent filesets. Information on how to program and use Starbase graphics is in the Starbase documentation.
HP Windows/9000 programming	DWINDOW MWINDOW PWINDOW	HP Windows/9000 demos. HP Windows/9000 on-line manual pages. HP Windows/9000 libraries for program development. Automatically loads dependent filesets. Information on how to program and use HP Windows/9000 is in the HP Windows/9000 documentation.
System accounting	PACCT	Automatically loads dependent filesets. Procedures on how to install and use system accounting are in the chapter "System Accounting" in the <i>System Administrator Manual</i> .
Run CE utilities	PCEUTIL	
Use source code control (SCCS).	PSCCS	Automatically loads dependent filesets.

**Table C-3. Programming Environment Functions  
Continued**

<b>Task</b>	<b>Fileset(s)</b>	<b>Notes</b>
Use DOS utilities	PAPPLIC	Commands used for prelocalization and localization of software.
Native Language Support	NCORE	
Configure and run an HP-UX cluster	PDSKLESS	
Write device drivers	PDRIVERS	

### **Filesets, Dependencies, and Sizes**

Before loading a fileset you must check to see that you have enough space on your file system. Check available space by typing:

```
df -t
```

You will see output similar to the following:

```

/      (/dev/dsk/0s0  ):      93918 blocks           135875 i-nodes
                          715904 total blocks       163840 total i-nodes
                          550394 used blocks        27965 used i-nodes
                          10 percent minfree

```

The size is given in 512-byte blocks. The available space in the example above is 93 918 blocks. Now you must use tables Figure C-2 and Figure C-3 to determine the space required for the new fileset. These tables list the fileset, its partition, its dependencies, and its size. *Note that the size is in 512-byte blocks.* For each fileset, you must take the sum of the fileset's size and all its dependencies' sizes. You do not need to add any dependencies that are already loaded on your system (look in the `/etc/filesets` directory to see if the fileset is already loaded).

If you know the specific command you need on your system, but can't decide which fileset it is in from the tables that follow, use the procedures in the section "How to List Files" to find the command you want. Once you know what fileset the file is in, use the `/etc/update` command to load the fileset onto your system. Note, however, that you cannot list the files contained on your installation or update media until you have installed AXE.

Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
KERNEL	1756	SYS_CORE	ACORE	9113	SYS_CORE
			CORE	868	SYS_CORE
			TOOL	561	SYS_CORE
			SSIT	456	SYS_CORE
			ACAT	5	SYS_CORE
			APAM	340	SYS_CORE
AMANUAL	1139	TEXT	ATEXT	71	TEXT
ASTARBAS	3961	STARBASE	AFA_FM	304	STARBASE
			AGRM	78	X11WINDOWS
AFA_FM	304	STARBASE	none		
AX11SERV	3208	X11WINDOWS	AGRM	78	X11WINDOWS
AX11RENV	8237	X11WINDOWS	AX11SERV	3208	X11WINDOWS
			AGRM	78	X11WINDOWS
AX11MAN	346	X11WINDOWS	ATEXT	71	TEXT
AWINDOW	2558	WINDOW	AFA_FM	304	STARBASE
			ASTARBAS	3961	STARBASE
			ATEXT	71	TEXT
			AGRM	78	X11WINDOWS
ACMD	1721	SYS_TOOLS	none		
ABCMD	2722	SYS_TOOLS	none		
AC	1443	PROG_LANGS	APROG	697	PROG_LANGS
AFPA	53	PROG_LANGS	none		
ATRMINFO	535	MISC_UTILS	none		
ACONFIG	3757	MISC_UTILS	AC	1443	PROG_LANGS
			APROG	697	PROG_LANGS
			LAN	1700	MISC_UTILS
AAUDIT	628	MISC_UTILS	none		

### AXE Filesets and Their Dependencies

Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
NAMERICA	10	NLS	none		
NCFRENCH	10	NLS	none		
NCHINESS	10	NLS	none		
NCHINEST	10	NLS	none		
NDANISH	10	NLS	none		
NDUTCH	10	NLS	none		
NENGLISH	10	NLS	none		
NFINNISH	10	NLS	none		
NFRENCH	10	NLS	none		
NGERMAN	10	NLS	none		
NGREEK	10	NLS	none		
NICELAND	10	NLS	none		
NITALIAN	10	NLS	none		
NJAPANES	20	NLS	none		
NKOREAN	10	NLS	none		
NNORWEGI	10	NLS	none		
NPORTUGU	10	NLS	none		
NSPANISH	10	NLS	none		
NSWEDISH	10	NLS	none		

**Figure C-2. AXE Filesets and Their Dependencies**

Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
PCORE	3821	SYS_CORE	FTN_CORE	492	SYS_CORE
			PAS_CORE	588	SYS_CORE
			ACMD	1721	SYS_TOOLS
PDSKLESS	394	SYS_CORE	none		
PTEXT	2413	TEXT	ATEXT	71	TEXT
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
HPUX_MAN	4050	TEXT	ATEXT	71	TEXT
			PTEXT	2413	TEXT
			PFILTER	909	SYS_TOOLS
PSTARBAS	6752	STARBASE	ACMD	1721	SYS_TOOLS
			AFA_FM	304	STARBASE
			ASTARBAS	3961	STARBASE
DSTARBAS	2146	STARBASE	AC	1443	PROG_LANGS
			APROG	697	PROG_LANGS
			PC	3870	PROG_LANGS
			PPROG	4696	PROG_LANGS
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
			AGRM	78	X11WINDOWS
			AFA_FM	304	STARBASE
			PSTARBAS	6752	STARBASE
			ASTARBAS	3961	STARBASE
MSTARBAS	616	STARBASE	APROG	697	PROG_LANGS
			AC	1443	PROG_LANGS
			PC	3870	PROG_LANGS
			PPROG	4696	PROG_LANGS
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
			AGRM	78	X11WINDOWS
			PTEXT	2413	TEXT
			ATEXT	71	TEXT
			PFILTER	909	SYS_TOOLS
ACMD	1721	SYS_TOOLS			

### PE Filesets and Their Dependencies



Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
PFA_FM	214	STARBASE	AFA_FM	304	STARBASE
			APROG	697	PROG_LANGS
			PPROG	4696	PROG_LANGS
			AC	1443	PROG_LANGS
			PC	3870	PROG_LANGS
			ATEXT	71	TEXT
			PTEXT	2413	TEXT
			ACMD	1721	SYS_TOOLS
			PFILTER	909	SYS_TOOLS
			PX11PENV	4622	X11WINDOWS
AX11RENV	8237	X11WINDOWS			
AGRM	78	X11WINDOWS			
PX11MAN	1228	X11WINDOWS	ATEXT	71	TEXT
			PTEXT	2413	TEXT
			PFILTER	909	SYS_TOOLS
PWINDOW	190	WINDOW	ACMD	1721	SYS_TOOLS
			AWINDOW	2558	WINDOW
			AFA_FM	304	STARBASE
			ASTARBAS	3961	STARBASE
DWINDOW	642	WINDOW	ATEXT	71	TEXT
			AGRM	78	X11WINDOWS
			AWINDOW	2558	WINDOW
			PWINDOW	190	WINDOW
			ATEXT	71	TEXT
			PTEXT	2413	TEXT
			AFA_FM	304	STARBASE
			PFA_FM	214	STARBASE
			ASTARBAS	3961	STARBASE
			PSTARBAS	6752	STARBASE
			AGRM	78	X11WINDOWS
			AC	1443	PROG_LANGS
			APROG	697	PROG_LANGS
			PC	3870	PROG_LANGS
			PPROG	4696	PROG_LANGS
PFILTER	909	SYS_TOOLS			
ACMD	1721	SYS_TOOLS			
PAPPLIC	2488	MISC_UTILS			

### PE Filesets and their Dependencies

Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
MWINDOW	406	WINDOW	PTEXT	2413	TEXT
			ATEXT	71	TEXT
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
PRLP	584	SYS_TOOLS	none		
PSYSCOM	4854	SYS_TOOLS	ACMD	1721	SYS_TOOLS
PFILTER	909	SYS_TOOLS	ACMD	1721	SYS_TOOLS
PCEUTIL	5493	SYS_TOOLS	ACMD	1721	SYS_TOOLS
			AC	1443	PROG_LANGS
			APROG	697	PROG_LANGS
PDRIVERS	537	SYS_TOOLS	none		
PPROG	4696	PROG_LANGS	APROG	697	PROG_LANGS
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
			AC	1443	PROG_LANGS
PC	3870	PROG_LANGS	PPROG	4696	PROG_LANGS
			APROG	697	PROG_LANGS
			PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
PFORTRAN	92	PROG_LANGS	none		
PAPPLIC	2488	MISC_UTILS	PFILTER	909	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
PPERIPH	330	MISC_UTILS	none		
PSCCS	2849	MISC_UTILS	ACMD	1721	SYS_TOOLS
			PSYSCOM	4854	SYS_TOOLS
			ACMD	1721	SYS_TOOLS
PACCT	1196	MISC_UTILS	PFILTER	909	SYS_TOOLS
			PFILTER	909	SYS_TOOLS
NCORE	994	NLS	none		
NNCOMPUT	277	NLS	NCORE	994	NLS

### PE Filesets and their Dependencies

Fileset	Size in Blocks	Partition	Dependent Filesets		
			Name	Size	Partition
PORT-NLS	14	NLS	NAMERICA	10	NLS
			NCFRENCH	10	NLS
			NCHINESS	10	NLS
			NCHINEST	10	NLS
			NDANISH	10	NLS
			NDUTCH	10	NLS
			NENGLISH	10	NLS
			NFINNISH	10	NLS
			NFRENCH	10	NLS
			NGERMAN	10	NLS
			NGREEK	10	NLS
			NICELAND	10	NLS
			NITALIAN	10	NLS
			NJAPANES	20	NLS
			NKOREAN	10	NLS
			NNORWEGI	10	NLS
			NPORTUGU	10	NLS
			NSPANISH	10	NLS
			NSWEDISH	10	NLS

**Figure C-3. PE Filesets and their Dependencies**

---

## Recovering Lost Files

If any of HP-UX commands are removed or somehow destroyed, you can use `/etc/update` to re-load the fileset containing the commands. Before you can use `/etc/update` check to see that your system has the following files (required for the `/etc/update` process to run):

```

/bin/cpio
/bin/mkdir
/bin/pwd
/usr/bin/lifcp
/usr/bin/tcio
/etc/update
/etc/sysrm

```

If these files are not on the system, you must use the procedures in the following section to copy them back onto the system. All the files required for `/etc/update` are in the `SYS_CORE` partition.

## Recovering Files Required For Update

### If you installed from a cartridge tape:

1. Make sure the tape is not write protected. The arrow on the tape should not be pointing to the word "SAFE".
2. Insert tape in drive; wait for the drive light to remain off.
3. Mount the tape to an empty directory by typing:

```
mount /dev/ct/s0 /disc
```

where `/dev/ct/s0` is a block device file.

4. Using the following command lines, copy all the files you need onto your hard disk:

```
cp /disc/bin/cpio /bin
cp /disc/bin/mkdir /bin
cp /disc/bin/pwd /bin
cp /disc/usr/bin/lifcp /usr/bin
cp /disc/urs/bin/tcio /usr/bin
cp /disc/etc/update /etc
cp /disc/etc/sysrm /etc
```

5. Type: `sync`
6. Unmount the cartridge tape by typing:

```
umount /dev/ct/s0
```

### If you installed from micro disks:

1. Make sure micro disk 2 (grey label) is not write protected.
2. Insert micro disk 2 (grey label) into the drive.
3. Mount the micro disk to an empty directory by typing:

```
mount /dev/dsk/1s0 /disc
```

where `/dev/dsk/1s0` must be a block device.

4. Copy the required commands from flexible disk 2 onto your hard disk by typing:

```
cp /disc/bin/cpio /bin
cp /disc/bin/mkdir /bin
cp /disc/bin/pwd /bin
cp /disc/usr/bin/lifcp /usr/bin
cp /disc/usr/bin/tcio /usr/bin
cp /disc/etc/update /etc
cp /disc/etc/sysrm /etc
```

5. Unmount the second micro disk by typing:

```
umount /dev/dsk/1s0
```

## How to List Files

If you are looking for a specific file or command you may need to list the files on your installation tape or flexible disks. You will need to have the LIF utilities loaded onto your system to do this. Use the procedures in the “Installing Optional Software and Updating Your HP-UX Kernel” in this manual to load the LIF utilities onto your system. The LIF utilities are in the ABCMD fileset.

To list the files do the following:

1. Insert the flexible disk or tape into the drive. If you are using cartridge tape wait for the light on the tape drive to remain off.
2. Get a list of archive file names by typing:

```
lifls -l /dev/rct/c0
```

This will show the names of the filesets on the media. Each name will begin with an extra “A”.

3. Get a list of each file or command in an archive file by typing:

```
lifcp -r /dev/rct/c0:lif_archive_name - | cpio -it
```

where *lif\_archive\_name* is the name of one of the archive files listed by executing the command in step 2. You must re-execute the `lifcp` command for each archive file listed from step 2.

---

## List of AXE System Files

The following is a list of the AXE partitions and filesets, and the files contained within them. This list may vary slightly from what is actually provided on the installation media. Also, some of the files are not actually on the media but are created by the customization scripts. If you want a particular file, load the whole fileset.

### Partition SYS\_CORE

#### Fileset ACAT

/system/ACAT/revlist	/usr/lib/ioerrs
/usr/lib/escerrs	/usr/lib/syserrs

#### Fileset ACORE

/bin/basename	/etc/conf/.clustr.help
/bin/cat	/etc/conf/.dcnode.help
/bin/chacl	/etc/conf/.dev.help
/bin/chgrp	/etc/conf/.dfile.help
/bin/chmod	/etc/conf/.dkless.help
/bin/chown	/etc/conf/.duser.help
/bin/cnodes	/etc/conf/.exit.help
/bin/cp	/etc/conf/.int.help
/bin/date	/etc/conf/.intro.help
/bin/dd	/etc/conf/.lp.help
/bin/df	/etc/conf/.prod.help
/bin/dirname	/etc/conf/.term.help
/bin/echo	/etc/conf/.user.help
/bin/expr	/etc/conf/.wkstn.help
/bin/false	/etc/convertfs
/bin/find	/etc/copyright
/bin/getcontext	/etc/devnm

/bin/grep	/etc/diskdescribe
/bin/grget	/etc/fbackup
/bin/hostname	/etc/fbackupprdr
/bin/hp9000s200	/etc/fbackupwrtr
/bin/hp9000s300	/etc/frecover
/bin/hp9000s500	/etc/fsck
/bin/hp9000s800	/etc/fsclean
/bin/ipcrm	/etc/getty
/bin/ipcs	/etc/init
/bin/kill	/etc/killall
/bin/l	/etc/link
/bin/line	/etc/mkfs
/bin/ll	/etc/mknod
/bin/ln	/etc/mount
/bin/login	/etc/mvdir
/bin/ls	/etc/newconfig/.profile
/bin/lsacl	/etc/newconfig/backup
/bin/lsf	/etc/newconfig/backupf
/bin/lsr	/etc/newconfig/bcheckrc
/bin/lsx	/etc/newconfig/brc
/bin/mv	/etc/newconfig/checklist
/bin/nohup	/etc/newconfig/d.login
/bin/passwd	/etc/newconfig/d.profile
/bin/pdp11	/etc/newconfig/disktab
/bin/ps	/etc/newconfig/gettydefs
/bin/pwget	/etc/newconfig/group
/bin/rm	/etc/newconfig/inittab
/bin/rmdir	/etc/newconfig/lpschedCDF
/bin/rsh	/etc/newconfig/motd
/bin/sed	/etc/newconfig/multi_user
/bin/sleep	/etc/newconfig/passwd
/bin/sort	/etc/newconfig/profile
/bin/stty	/etc/newconfig/pstatus
/bin/su	/etc/newconfig/qstatus
/bin/sync	/etc/newconfig/rc
/bin/touch	/etc/newconfig/rebootCDF
/bin/true	/etc/newconfig/slpCDF
/bin/u370	/etc/newconfig/ttytype
/bin/u3b	/etc/newconfig/tztab
/bin/u3b10	/etc/newfs
/bin/u3b2	/etc/reboot
/bin/u3b5	/etc/reconfig
/bin/uname	/etc/setmnt
/bin/vax	/etc/shutdown

/bin/wc	/etc/single_user
/bin/who	/etc/swapon
/etc/boot	/etc/syncer
/etc/conf/.acnode.help	/etc/telinit
/etc/conf/.auser.help	/etc/unlink
/etc/wall	/usr/lib/terminfo/2/2393a
/system/ACORE/ACORE.links	/usr/lib/terminfo/2/2394
/system/ACORE/CDFcustomize	/usr/lib/terminfo/2/2394A
/system/ACORE/customize	/usr/lib/terminfo/2/2394a
/system/ACORE/revlist	/usr/lib/terminfo/2/2397
/users/guest/.profile	/usr/lib/terminfo/2/2397A
/usr/bin/asa	/usr/lib/terminfo/2/2397a
/usr/bin/awk	/usr/lib/terminfo/2/2621
/usr/bin/banner	/usr/lib/terminfo/2/2621-48
/usr/bin/cancel	/usr/lib/terminfo/2/2621-ba
/usr/bin/cut	/usr/lib/terminfo/2/2621-fl
/usr/bin/disable	/usr/lib/terminfo/2/2621-nl
/usr/bin/enable	/usr/lib/terminfo/2/2621-nt
/usr/bin/fstomnt	/usr/lib/terminfo/2/2621-pb
/usr/bin/getaccess	/usr/lib/terminfo/2/2621-wl
/usr/bin/lp	/usr/lib/terminfo/2/2621A
/usr/bin/lpr	/usr/lib/terminfo/2/2621a
/usr/bin/lpstat	/usr/lib/terminfo/2/2621k45
/usr/bin/mediainit	/usr/lib/terminfo/2/2621nl
/usr/bin/more	/usr/lib/terminfo/2/2621nt
/usr/bin/page	/usr/lib/terminfo/2/2621wl
/usr/bin/reset	/usr/lib/terminfo/2/2622
/usr/bin/slp	/usr/lib/terminfo/2/2622a
/usr/bin/tabs	/usr/lib/terminfo/2/2622p
/usr/bin/tar	/usr/lib/terminfo/2/2623
/usr/bin/tset	/usr/lib/terminfo/2/2623a
/usr/bin/what	/usr/lib/terminfo/2/2623p
/usr/lib/accept	/usr/lib/terminfo/2/2624
/usr/lib/expreserve	/usr/lib/terminfo/2/2624a
/usr/lib/hp2686a	/usr/lib/terminfo/2/2624p
/usr/lib/hp2934a	/usr/lib/terminfo/2/2625
/usr/lib/lpadmin	/usr/lib/terminfo/2/2626
/usr/lib/lpfx	/usr/lib/terminfo/2/2626-12
/usr/lib/lpmove	/usr/lib/terminfo/2/2626-12-s
/usr/lib/lprpp	/usr/lib/terminfo/2/2626-12x40
/usr/lib/lpsched	/usr/lib/terminfo/2/2626-ns
/usr/lib/lpshut	/usr/lib/terminfo/2/2626-s
/usr/lib/more.help	/usr/lib/terminfo/2/2626-x40
/usr/lib/nls/config	/usr/lib/terminfo/2/2626A
/usr/lib/reject	/usr/lib/terminfo/2/2626P



/usr/lib/reverse	/usr/lib/terminfo/2/2626a
/usr/lib/tabset/3101	/usr/lib/terminfo/2/2626p
/usr/lib/tabset/9826	/usr/lib/terminfo/2/2627
/usr/lib/tabset/9836	/usr/lib/terminfo/2/2627a
/usr/lib/tabset/9837	/usr/lib/terminfo/2/2627p
/usr/lib/tabset/98700	/usr/lib/terminfo/2/2628
/usr/lib/tabset/bee hive	/usr/lib/terminfo/2/262x
/usr/lib/tabset/diablo	/usr/lib/terminfo/2/2640
/usr/lib/tabset/std	/usr/lib/terminfo/2/2640a
/usr/lib/tabset/stdcrt	/usr/lib/terminfo/2/2640b
/usr/lib/tabset/teleray	/usr/lib/terminfo/2/2644
/usr/lib/tabset/vt100	/usr/lib/terminfo/2/2645
/usr/lib/tabset/xerox1720	/usr/lib/terminfo/2/2647
/usr/lib/terminfo/1/1003	/usr/lib/terminfo/2/2647F
/usr/lib/terminfo/1/1003a	/usr/lib/terminfo/2/2648
/usr/lib/terminfo/1/1003g	/usr/lib/terminfo/2/2648A
/usr/lib/terminfo/1/110	/usr/lib/terminfo/2/2648a
/usr/lib/terminfo/1/150	/usr/lib/terminfo/2/2703
/usr/lib/terminfo/1/150A	/usr/lib/terminfo/2/2709
/usr/lib/terminfo/1/150a	/usr/lib/terminfo/3/300h
/usr/lib/terminfo/2/2382	/usr/lib/terminfo/3/300I
/usr/lib/terminfo/2/2392	/usr/lib/terminfo/7/700-41
/usr/lib/terminfo/2/2392A	/usr/lib/terminfo/9/9020
/usr/lib/terminfo/2/2392a	/usr/lib/terminfo/9/9816te
/usr/lib/terminfo/2/2393	/usr/lib/terminfo/9/9816teb
/usr/lib/terminfo/2/2393A	/usr/lib/terminfo/9/98204b
/usr/lib/terminfo/9/9826	/usr/lib/terminfo/h/hp2397a
/usr/lib/terminfo/9/9826ite	/usr/lib/terminfo/h/hp2621
/usr/lib/terminfo/9/9836	/usr/lib/terminfo/h/hp2621-ba
/usr/lib/terminfo/9/9836a	/usr/lib/terminfo/h/hp2621-fl
/usr/lib/terminfo/9/9836ite	/usr/lib/terminfo/h/hp2621-nl
/usr/lib/terminfo/9/9836te	/usr/lib/terminfo/h/hp2621-nt
/usr/lib/terminfo/9/9836teb	/usr/lib/terminfo/h/hp2621-pb
/usr/lib/terminfo/9/9837	/usr/lib/terminfo/h/hp2621-wl
/usr/lib/terminfo/9/9837a	/usr/lib/terminfo/h/hp2621A
/usr/lib/terminfo/9/9837ite	/usr/lib/terminfo/h/hp2621a
/usr/lib/terminfo/9/98541	/usr/lib/terminfo/h/hp2621k45
/usr/lib/terminfo/9/98541a	/usr/lib/terminfo/h/hp2621nl
/usr/lib/terminfo/9/98542	/usr/lib/terminfo/h/hp2621nt
/usr/lib/terminfo/9/98542a	/usr/lib/terminfo/h/hp2621wl
/usr/lib/terminfo/9/98543	/usr/lib/terminfo/h/hp2622
/usr/lib/terminfo/9/98543a	/usr/lib/terminfo/h/hp2622a
/usr/lib/terminfo/9/98544	/usr/lib/terminfo/h/hp2622p
/usr/lib/terminfo/9/98544a	/usr/lib/terminfo/h/hp2623
/usr/lib/terminfo/9/98545	/usr/lib/terminfo/h/hp2623a

/usr/lib/terminfo/9/98545a	/usr/lib/terminfo/h/hp2623p
/usr/lib/terminfo/9/98546	/usr/lib/terminfo/h/hp2624
/usr/lib/terminfo/9/98546a	/usr/lib/terminfo/h/hp2624a
/usr/lib/terminfo/9/98547	/usr/lib/terminfo/h/hp2624p
/usr/lib/terminfo/9/98547a	/usr/lib/terminfo/h/hp2625
/usr/lib/terminfo/9/98548	/usr/lib/terminfo/h/hp2626
/usr/lib/terminfo/9/98548a	/usr/lib/terminfo/h/hp2626a
/usr/lib/terminfo/9/98549	/usr/lib/terminfo/h/hp2626p
/usr/lib/terminfo/9/98549a	/usr/lib/terminfo/h/hp2627
/usr/lib/terminfo/9/98550	/usr/lib/terminfo/h/hp2627a
/usr/lib/terminfo/9/98550a	/usr/lib/terminfo/h/hp2627p
/usr/lib/terminfo/9/98700	/usr/lib/terminfo/h/hp2628
/usr/lib/terminfo/9/98700a	/usr/lib/terminfo/h/hp2640a
/usr/lib/terminfo/9/98710	/usr/lib/terminfo/h/hp2640b
/usr/lib/terminfo/9/98710a	/usr/lib/terminfo/h/hp2644
/usr/lib/terminfo/9/98720	/usr/lib/terminfo/h/hp2645
/usr/lib/terminfo/9/98720a	/usr/lib/terminfo/h/hp2648
/usr/lib/terminfo/9/98721	/usr/lib/terminfo/h/hp2648a
/usr/lib/terminfo/9/98721a	/usr/lib/terminfo/h/hp2703
/usr/lib/terminfo/9/98730	/usr/lib/terminfo/h/hp2709
/usr/lib/terminfo/9/98730a	/usr/lib/terminfo/h/hp300h
/usr/lib/terminfo/9/98731	/usr/lib/terminfo/h/hp300l
/usr/lib/terminfo/9/98731a	/usr/lib/terminfo/h/hp44
/usr/lib/terminfo/X/X-hpterm	/usr/lib/terminfo/h/hp45
/usr/lib/terminfo/a/ansi	/usr/lib/terminfo/h/hp9020
/usr/lib/terminfo/c/c1003	/usr/lib/terminfo/h/hp9816te
/usr/lib/terminfo/c/c1003a	/usr/lib/terminfo/h/hp9816teb
/usr/lib/terminfo/c/c1003g	/usr/lib/terminfo/h/hp98204b
/usr/lib/terminfo/d/dumb	/usr/lib/terminfo/h/hp9826
/usr/lib/terminfo/h/hp	/usr/lib/terminfo/h/hp9836
/usr/lib/terminfo/h/hp110	/usr/lib/terminfo/h/hp9836a
/usr/lib/terminfo/h/hp150	/usr/lib/terminfo/h/hp9836te
/usr/lib/terminfo/h/hp150A	/usr/lib/terminfo/h/hp9836teb
/usr/lib/terminfo/h/hp150a	/usr/lib/terminfo/h/hp9837
/usr/lib/terminfo/h/hp2382	/usr/lib/terminfo/h/hp9837a
/usr/lib/terminfo/h/hp2382a	/usr/lib/terminfo/h/hp98541
/usr/lib/terminfo/h/hp2392	/usr/lib/terminfo/h/hp98541a
/usr/lib/terminfo/h/hp2392A	/usr/lib/terminfo/h/hp98542
/usr/lib/terminfo/h/hp2392a	/usr/lib/terminfo/h/hp98542a
/usr/lib/terminfo/h/hp2393	/usr/lib/terminfo/h/hp98543
/usr/lib/terminfo/h/hp2393A	/usr/lib/terminfo/h/hp98543a
/usr/lib/terminfo/h/hp2393a	/usr/lib/terminfo/h/hp98544
/usr/lib/terminfo/h/hp2394	/usr/lib/terminfo/h/hp98544a
/usr/lib/terminfo/h/hp2394A	/usr/lib/terminfo/h/hp98545
/usr/lib/terminfo/h/hp2394a	/usr/lib/terminfo/h/hp98545a

/usr/lib/terminfo/h/hp2397	/usr/lib/terminfo/h/hp98546
/usr/lib/terminfo/h/hp2397A	/usr/lib/terminfo/h/hp98546a
/usr/lib/terminfo/h/hp98547	/usr/lib/terminfo/h/hpsub
/usr/lib/terminfo/h/hp98547a	/usr/lib/terminfo/h/hpterm
/usr/lib/terminfo/h/hp98548	/usr/lib/terminfo/k/k45
/usr/lib/terminfo/h/hp98548a	/usr/lib/terminfo/n/nomad
/usr/lib/terminfo/h/hp98549	/usr/lib/terminfo/u/unknown
/usr/lib/terminfo/h/hp98549a	/usr/spool/lp/model/PCL1
/usr/lib/terminfo/h/hp98550	/usr/spool/lp/model/PCL2
/usr/lib/terminfo/h/hp98550a	/usr/spool/lp/model/PCL3
/usr/lib/terminfo/h/hp98700	/usr/spool/lp/model/deskjet
/usr/lib/terminfo/h/hp98700a	/usr/spool/lp/model/dumb
/usr/lib/terminfo/h/hp98710	/usr/spool/lp/model/hp2225a
/usr/lib/terminfo/h/hp98710a	/usr/spool/lp/model/hp2227a
/usr/lib/terminfo/h/hp98720	/usr/spool/lp/model/hp2560
/usr/lib/terminfo/h/hp98720a	/usr/spool/lp/model/hp2631g
/usr/lib/terminfo/h/hp98721	/usr/spool/lp/model/hp2686a
/usr/lib/terminfo/h/hp98721a	/usr/spool/lp/model/hp2934a
/usr/lib/terminfo/h/hp98730	/usr/spool/lp/model/hp3630a
/usr/lib/terminfo/h/hp98730a	/usr/spool/lp/model/laserjet
/usr/lib/terminfo/h/hp98731	/usr/spool/lp/model/quietjet
/usr/lib/terminfo/h/hp98731a	/usr/spool/lp/model/ruggedwriter
/usr/lib/terminfo/h/hpex	/usr/spool/lp/model/thinkjet

## Fileset APAM

/bin/pam	/system/APAM/customize
/etc/newconfig/d.Autost	/system/APAM/revlist
/etc/newconfig/d.environ	/usr/lib/nls/n-computer/pam.cat

## Fileset CORE

/bin/cpio	/system/CORE/revlist
/bin/mkdir	/usr/bin/compress
/bin/pwd	/usr/bin/lifcp
/bin/sh	/usr/bin/tcio
/etc/umount	/usr/bin/uncompress
/system/CORE/customize	/usr/bin/zcat

## Fileset KERNEL

/hp-ux	/system/KERNEL/customize
/system/KERNEL/CDFcustomize	/system/KERNEL/revlist

## Fileset SSIT

/etc/newconfig/README	/system/SSIT/convert
/etc/newconfig/Update_info/to.6.5	/system/SSIT/getcnodeid
/system/SSIT/CDFcustomize	/system/SSIT/getservid
/system/SSIT/cleanup	/system/SSIT/revlist

## Fileset TOOL

/etc/sysrm	/system/TOOL/rebuild.ckerns
/etc/update	/system/TOOL/revlist
/system/TOOL/cpio	

## Partition TEXT

### Fileset AMANUAL

/system/AMANUAL/customize	/usr/man/cat1.Z/lpstat.1
/usr/man/cat1.Z/asa.1	/usr/man/cat1.Z/ls.1
/usr/man/cat1.Z/awk.1	/usr/man/cat1.Z/lsacl.1
/usr/man/cat1.Z/banner.1	/usr/man/cat1.Z/lsf.1
/usr/man/cat1.Z/basename.1	/usr/man/cat1.Z/lxr.1
/usr/man/cat1.Z/bfs.1	/usr/man/cat1.Z/lxx.1
/usr/man/cat1.Z/bifchgrp.1	/usr/man/cat1.Z/man.1
/usr/man/cat1.Z/bifchmod.1	/usr/man/cat1.Z/mediainit.1
/usr/man/cat1.Z/bifchown.1	/usr/man/cat1.Z/mkdir.1
/usr/man/cat1.Z/bifcpl.1	/usr/man/cat1.Z/more.1
/usr/man/cat1.Z/biffind.1	/usr/man/cat1.Z/mv.1
/usr/man/cat1.Z/biflfs.1	/usr/man/cat1.Z/nice.1
/usr/man/cat1.Z/bifmkdir.1	/usr/man/cat1.Z/nohup.1
/usr/man/cat1.Z/bifrm.1	/usr/man/cat1.Z/page.1
/usr/man/cat1.Z/bifrmkdir.1	/usr/man/cat1.Z/pam.1
/usr/man/cat1.Z/cancel.1	/usr/man/cat1.Z/passwd.1
/usr/man/cat1.Z/cat.1	/usr/man/cat1.Z/pr.1
/usr/man/cat1.Z/cd.1	/usr/man/cat1.Z/ps.1
/usr/man/cat1.Z/chacl.1	/usr/man/cat1.Z/pwd.1
/usr/man/cat1.Z/chgrp.1	/usr/man/cat1.Z/pwget.1
/usr/man/cat1.Z/chmod.1	/usr/man/cat1.Z/red.1
/usr/man/cat1.Z/chown.1	/usr/man/cat1.Z/rm.1
/usr/man/cat1.Z/cmp.1	/usr/man/cat1.Z/rmdir.1
/usr/man/cat1.Z/cnodes.1	/usr/man/cat1.Z/rsh.1
/usr/man/cat1.Z/compress.1	/usr/man/cat1.Z/sed.1

```

/usr/man/cat1.Z/cp.1
/usr/man/cat1.Z/cpio.1
/usr/man/cat1.Z/csh.1
/usr/man/cat1.Z/cut.1
/usr/man/cat1.Z/date.1
/usr/man/cat1.Z/dd.1
/usr/man/cat1.Z/dirname.1
/usr/man/cat1.Z/du.1
/usr/man/cat1.Z/echo.1
/usr/man/cat1.Z/ed.1
/usr/man/cat1.Z/edit.1
/usr/man/cat1.Z/egrep.1
/usr/man/cat1.Z/enable.1
/usr/man/cat1.Z/ex.1
/usr/man/cat1.Z/expr.1
/usr/man/cat1.Z/false.1
/usr/man/cat1.Z/fgrep.1
/usr/man/cat1.Z/find.1
/usr/man/cat1.Z/fold.1
/usr/man/cat1.Z/getaccess.1
/usr/man/cat1.Z/getcontext.1
/usr/man/cat1.Z/grep.1
/usr/man/cat1.Z/help.1
/usr/man/cat1.Z/hostname.1
/usr/man/cat1.Z/hp.1
/usr/man/cat1.Z/intro.1
/usr/man/cat1.Z/ipcrm.1
/usr/man/cat1.Z/ipcs.1
/usr/man/cat1.Z/kill.1
/usr/man/cat1.Z/ksh.1
/usr/man/cat1.Z/l.1
/usr/man/cat1.Z/lifcp.1
/usr/man/cat1.Z/lifinit.1
/usr/man/cat1.Z/lifls.1
/usr/man/cat1.Z/lifrename.1
/usr/man/cat1.Z/lifrm.1
/usr/man/cat1.Z/line.1
/usr/man/cat1.Z/ll.1
/usr/man/cat1.Z/ln.1
/usr/man/cat1.Z/login.1
/usr/man/cat1.Z/lp.1
/usr/man/cat1m.Z/fbackup.1m
/usr/man/cat1m.Z/frecover.1m
/usr/man/cat1m.Z/fsck.1m
/usr/man/cat1m.Z/fsclean.1m
/usr/man/cat1.Z/sh.1
/usr/man/cat1.Z/sleep.1
/usr/man/cat1.Z/slp.1
/usr/man/cat1.Z/sort.1
/usr/man/cat1.Z/stty.1
/usr/man/cat1.Z/su.1
/usr/man/cat1.Z/tabs.1
/usr/man/cat1.Z/tail.1
/usr/man/cat1.Z/tar.1
/usr/man/cat1.Z/tcio.1
/usr/man/cat1.Z/tee.1
/usr/man/cat1.Z/touch.1
/usr/man/cat1.Z/tr.1
/usr/man/cat1.Z/true.1
/usr/man/cat1.Z/tset.1
/usr/man/cat1.Z/umask.1
/usr/man/cat1.Z/uname.1
/usr/man/cat1.Z/uncompress.1
/usr/man/cat1.Z/vi.1
/usr/man/cat1.Z/wc.1
/usr/man/cat1.Z/what.1
/usr/man/cat1.Z/whereis.1
/usr/man/cat1.Z/which.1
/usr/man/cat1.Z/who.1
/usr/man/cat1.Z/whoami.1
/usr/man/cat1.Z/zcat.1
/usr/man/cat1m.Z/accept.1m
/usr/man/cat1m.Z/audevent.1m
/usr/man/cat1m.Z/audisp.1m
/usr/man/cat1m.Z/audomon.1m
/usr/man/cat1m.Z/audsys.1m
/usr/man/cat1m.Z/audusr.1m
/usr/man/cat1m.Z/backup.1m
/usr/man/cat1m.Z/bcheckrc.1m
/usr/man/cat1m.Z/biffsc.1m
/usr/man/cat1m.Z/bifmkfs.1m
/usr/man/cat1m.Z/brc.1m
/usr/man/cat1m.Z/chroot.1m
/usr/man/cat1m.Z/config.1m
/usr/man/cat1m.Z/devnm.1m
/usr/man/cat1m.Z/df.1m
/usr/man/cat1m.Z/reboot.1m
/usr/man/cat1m.Z/reconfig.1m
/usr/man/cat1m.Z/reject.1m
/usr/man/cat1m.Z/setmnt.1m

```

/usr/man/cat1m.Z/fstomnt.1m	/usr/man/cat1m.Z/shutdown.1m
/usr/man/cat1m.Z/getty.1m	/usr/man/cat1m.Z/swapon.1m
/usr/man/cat1m.Z/init.1m	/usr/man/cat1m.Z/sync.1m
/usr/man/cat1m.Z/intro.1m	/usr/man/cat1m.Z/syncer.1m
/usr/man/cat1m.Z/killall.1m	/usr/man/cat1m.Z/telinit.1m
/usr/man/cat1m.Z/link.1m	/usr/man/cat1m.Z/umount.1m
/usr/man/cat1m.Z/lpadmin.1m	/usr/man/cat1m.Z/unlink.1m
/usr/man/cat1m.Z/lpmove.1m	/usr/man/cat1m.Z/wall.1m
/usr/man/cat1m.Z/lpsched.1m	/usr/man/cat4.Z/audit.4
/usr/man/cat1m.Z/lpshut.1m	/usr/man/cat4.Z/bif.4
/usr/man/cat1m.Z/mkdev.1m	/usr/man/cat4.Z/checklist.4
/usr/man/cat1m.Z/mkfs.1m	/usr/man/cat4.Z/innitab.4
/usr/man/cat1m.Z/mknod.1m	/usr/man/cat4.Z/intro.4
/usr/man/cat1m.Z/mkrs.1m	/usr/man/cat4.Z/lif.4
/usr/man/cat1m.Z/mount.1m	/usr/man/cat4.Z/passwd.4
/usr/man/cat1m.Z/mvdir.1m	/usr/man/cat4.Z/profile.4
/usr/man/cat1m.Z/newfs.1m	/usr/man/cat4.Z/ttytype.4
/usr/man/cat1m.Z/powerfail.1m	/usr/man/cat4.Z/tztab.4
/usr/man/cat1m.Z/rc.1m	

## Fileset ATEXT

/system/ATEXT/customize	/usr/bin/uncompressdir
/system/ATEXT/revlist	/usr/bin/zcomp
/usr/bin/compressdir	/usr/bin/zdiff
/usr/bin/man	

## Partition SYS\_TOOLS

### Fileset ABCMD

/system/ABCMD/customize	/usr/bin/lifrename
/system/ABCMD/revlist	/usr/bin/lifrm
/usr/bin/bifchgrp	/usr/bin/sdfchgrp
/usr/bin/bifchmod	/usr/bin/sdfchmod
/usr/bin/bifchown	/usr/bin/sdfchown
/usr/bin/bifcp	/usr/bin/sdfcp
/usr/bin/biffind	/usr/bin/sdfdf
/usr/bin/biffsck	/usr/bin/sdffind
/usr/bin/biffsdb	/usr/bin/sdffsck
/usr/bin/bifll	/usr/bin/sdffsdb
/usr/bin/bifln	/usr/bin/sdffll
/usr/bin/bifls	/usr/bin/sdffln
/usr/bin/bifmkdir	/usr/bin/sdffls

/usr/bin/bifmkfs	/usr/bin/sdfmkdir
/usr/bin/bifmv	/usr/bin/sdfmv
/usr/bin/bifrm	/usr/bin/sdfrm
/usr/bin/bifrmkdir	/usr/bin/sdfrmkdir
/usr/bin/hp	/usr/bin/tr
/usr/bin/lifinit	/usr/bin/whereis
/usr/bin/lifls	/usr/bin/whoami

## Fileset ACMD

/bin/cmp	/usr/lib/builtins/exec
/bin/csh	/usr/lib/builtins/foreach
/bin/du	/usr/lib/builtins/glob
/bin/ed	/usr/lib/builtins/hashstat
/bin/ksh	/usr/lib/builtins/history
/bin/nice	/usr/lib/builtins/if
/bin/pr	/usr/lib/builtins/jobs
/bin/red	/usr/lib/builtins/kill
/bin/rksh	/usr/lib/builtins/login
/bin/tail	/usr/lib/builtins/logout
/bin/tee	/usr/lib/builtins/newgrp
/etc/chroot	/usr/lib/builtins/nice
/etc/newconfig/csh.login	/usr/lib/builtins/nohup
/etc/newconfig/d.cshrc	/usr/lib/builtins/notify
/etc/newconfig/d.exrc	/usr/lib/builtins/onintr
/etc/newconfig/mkdev	/usr/lib/builtins/popd
/system/ACMD/customize	/usr/lib/builtins/pushd
/system/ACMD/revlist	/usr/lib/builtins/rehash
/usr/bin/bfs	/usr/lib/builtins/repeat
/usr/bin/e	/usr/lib/builtins/set
/usr/bin/edit	/usr/lib/builtins/setenv
/usr/bin/ex	/usr/lib/builtins/source
/usr/bin/fold	/usr/lib/builtins/stop
/usr/bin/getopt	/usr/lib/builtins/switch
/usr/bin/help	/usr/lib/builtins/time
/usr/bin/vedit	/usr/lib/builtins/umask
/usr/bin/vi	/usr/lib/builtins/unalias
/usr/bin/view	/usr/lib/builtins/unhash
/usr/bin/which	/usr/lib/builtins/unset
/usr/lib/builtins/@	/usr/lib/builtins/unsetenv
/usr/lib/builtins/alias	/usr/lib/builtins/wait
/usr/lib/builtins/alloc	/usr/lib/builtins/while
/usr/lib/builtins/case	/usr/lib/exrecovery
/usr/lib/builtins/cd	/usr/lib/help/default
/usr/lib/builtins/chdir	/usr/lib/help/he

/usr/lib/builtins/dirs  
/usr/lib/builtins/eval

/usr/lib/help/lib/help2

## Partition STARBASE

### Fileset AFA\_FM

/system/AFA\_FM/customize  
/system/AFA\_FM/revlist  
/usr/lib/raster/10x20/kana.8K  
/usr/lib/raster/10x20/lp.8U  
/usr/lib/raster/10x20/lp.b.8U  
/usr/lib/raster/8x16/lp.8U  
/usr/lib/raster/12x20/cour.0U  
/usr/lib/raster/12x20/cour.b.0U  
/usr/lib/raster/L6x15/lp.8U  
/usr/lib/raster/18x30/math.0M  
/usr/lib/raster/18x30/pica.8U  
/usr/lib/raster/6x8/SNF/lp.8U.scf  
/usr/lib/raster/6x8/lp.8U  
/usr/lib/raster/6x8/lp.b.8I  
/usr/lib/raster/6x8/math.8M  
/usr/lib/raster/7x10/SNF/lp.8U.scf  
/usr/lib/raster/7x10/lp.8U  
/usr/lib/raster/8x16/SNF/kana.8K.scf  
/usr/lib/raster/18x30/SNF/pica.8U.scf  
/usr/lib/raster/18x30/SNF/math.0M.scf  
/usr/lib/raster/12x20/SNF/cour.b.0U.scf  
/usr/lib/raster/12x20/SNF/cour.0U.scf  
/usr/lib/raster/10x20/SNF/kana.8K.scf  
/usr/lib/raster/8x16/SNF/lp.b.8U.scf  
/usr/lib/raster/10x20/SNF/lp.8U.scf  
/usr/lib/raster/8x16/SNF/lp.i.8U.scf  
/usr/lib/raster/10x20/SNF/lp.b.8U.scf  
/usr/lib/raster/8x16/SNF/math.0M.scf  
/usr/lib/raster/dflt/a/h/SNF/katakana.scf  
/usr/lib/raster/6x8/SNF/lp.b.8I.scf  
/usr/lib/raster/dflt/a/v/SNF/katakana.scf  
/usr/lib/raster/6x8/SNF/math.8M.scf  
/usr/lib/raster/dflt/a/v/katakana  
/usr/lib/raster/8x16/SNF/linedraw.0L.sc  
/usr/lib/raster/8x16/SNF/lp.8U.scf  
/usr/lib/raster/8x16/kana.8K  
/usr/lib/raster/8x16/linedraw.0L  
/usr/lib/raster/8x16/lp.8U  
/usr/lib/raster/8x16/lp.i.8U  
/usr/lib/raster/8x16/math.0M  
/usr/lib/raster/L6x15/SNF/lp.8U.scf  
/usr/lib/raster/dflt/b/v/katakana  
/usr/lib/raster/dflt/a/h/katakana  
/usr/lib/raster/dflt/a/l/SNF/katakana.scf  
/usr/lib/raster/dflt/a/l/katakana  
/usr/lib/raster/dflt/b/h/SNF/katakana.scf  
/usr/lib/raster/dflt/b/h/katakana  
/usr/lib/raster/dflt/b/l/SNF/katakana.scf  
/usr/lib/raster/dflt/b/l/katakana  
/usr/lib/raster/dflt/b/v/SNF/katakana.scf



## Fileset ASTARBAS

/system/ASTARBAS/customize	/usr/lib/starbase/stroke/font_info/58
/system/ASTARBAS/revlist	/usr/lib/starbase/stroke/font_info/6
/usr/bin/pcltrans	/usr/lib/starbase/stroke/font_info/8
/usr/lib/nls/american/sb.cat	/usr/lib/starbase/stroke/hproman/1
/usr/lib/nls/n-computer/sb.cat	/usr/lib/starbase/stroke/hproman/2
/usr/lib/starbase/char_sets	/usr/lib/starbase/stroke/hproman/4
/usr/lib/starbase/defaults	/usr/lib/starbase/stroke/hproman/58
/usr/lib/starbase/stroke/hproman/6	/usr/lib/starbase/stroke/hproman/8
/usr/lib/starbase/hp98556/gpuall.x	/usr/lib/starbase/stroke/jisascii/1
/usr/lib/starbase/hp98710/te_data	/usr/lib/starbase/stroke/jisascii/2
/usr/lib/starbase/hp98710/te_mcode	/usr/lib/starbase/stroke/jisascii/4
/usr/lib/starbase/hp98720/te_data	/usr/lib/starbase/stroke/jisascii/6
/usr/lib/starbase/hp98720/te_mcode	/usr/lib/starbase/stroke/jisascii/8
/usr/lib/starbase/hp98730/te_data	/usr/lib/starbase/stroke/katakana/1
/usr/lib/starbase/hp98730/te_mcode	/usr/lib/starbase/stroke/katakana/2
/usr/lib/starbase/sb_daemon_1.0	/usr/lib/starbase/stroke/katakana/4
/usr/lib/starbase/sb_daemon_1.1	/usr/lib/starbase/stroke/katakana/6
/usr/lib/starbase/sb_daemon_2.0	/usr/lib/starbase/stroke/katakana/8
/usr/lib/starbase/sb_daemon_2.1	/usr/lib/starbase/stroke/markers
/usr/lib/starbase/sb_daemon_2.2	/usr/lib/starbase/stroke/usascii/1
/usr/lib/starbase/sb_daemon_2.3	/usr/lib/starbase/stroke/usascii/2
/usr/lib/starbase/sb_daemon_2.4	/usr/lib/starbase/stroke/usascii/4
/usr/lib/starbase/sb_daemon_3.0	/usr/lib/starbase/stroke/usascii/58
/usr/lib/starbase/stroke/usascii/6	/usr/lib/starbase/stroke/usascii/8
/usr/lib/starbase/errors/american/sb.errors	
/usr/lib/starbase/errors/n-computer/sb.errors	
/usr/lib/starbase/stroke/font_info/1	
/usr/lib/starbase/stroke/font_info/2	
/usr/lib/starbase/stroke/font_info/4	

## Partition X11WINDOWS

### Fileset AGRM

/system/AGRM/revlist	/usr/lib/grmd
----------------------	---------------

### Fileset AX11MAN

/system/AX11MAN/customize	/usr/man/cat1.Z/xinitcolorm.1
/system/AX11MAN/revlist	/usr/man/cat1.Z/xload.1
/usr/man/cat1.Z/X.1	/usr/man/cat1.Z/xmodmap.1
/usr/man/cat1.Z/Xserver.1	/usr/man/cat1.Z/xpr.1
/usr/man/cat1.Z/bitmap.1	/usr/man/cat1.Z/xprkbd.1

```

/usr/man/cat1.Z/gwindstop.1
/usr/man/cat1.Z/hpterm.1
/usr/man/cat1.Z/hpwm.1
/usr/man/cat1.Z/resize.1
/usr/man/cat1.Z/rgb.1
/usr/man/cat1.Z/sb2xwd.1
/usr/man/cat1.Z/uwm.1
/usr/man/cat1.Z/x11start.1
/usr/man/cat1.Z/xclock.1
/usr/man/cat1.Z/xfc.1
/usr/man/cat1.Z/xfd.1
/usr/man/cat1.Z/xhost.1
/usr/man/cat1.Z/xinit.1

/usr/man/cat1.Z/xrdb.1
/usr/man/cat1.Z/xrefresh.1
/usr/man/cat1.Z/xsethru.1
/usr/man/cat1.Z/xset.1
/usr/man/cat1.Z/xsetroot.1
/usr/man/cat1.Z/xterm.1
/usr/man/cat1.Z/xwcreate.1
/usr/man/cat1.Z/xwd.1
/usr/man/cat1.Z/xwd2sb.1
/usr/man/cat1.Z/xwdestroy.1
/usr/man/cat1.Z/xwininfo.1
/usr/man/cat1.Z/xwud.1

```

## Fileset AX11RENV

```

/etc/newconfig/Hpwm
/etc/newconfig/rgb.txt
/etc/newconfig/sys.Xdefaults
/etc/newconfig/sys.x11start
/etc/newconfig/system.hpwmrc
/etc/newconfig/system.uwmrc
/etc/newconfig/xcolormap
/system/AX11RENV/customize
/system/AX11RENV/revlist
/usr/bin/X11/bitmap
/usr/bin/X11/gwind
/usr/bin/X11/gwindstop
/usr/bin/X11/hpterm
/usr/bin/X11/hpwm
/usr/bin/X11/resize
/usr/bin/X11/rgb
/usr/bin/X11/sb2xwd
/usr/bin/X11/uwm
/usr/bin/X11/xclock
/usr/bin/X11/xfc
/usr/bin/X11/xfd
/usr/bin/X11/xhost
/usr/bin/X11/xinit
/usr/bin/X11/xinitcolormap
/usr/bin/X11/xload
/usr/bin/X11/xmodmap
/usr/bin/X11/xpr
/usr/bin/X11/xprkbd
/usr/bin/X11/xrdb

/usr/include/X11/bitmaps/1x1
/usr/include/X11/bitmaps/2x2
/usr/include/X11/bitmaps/black
/usr/include/X11/bitmaps/boxes
/usr/include/X11/bitmaps/cntr_ptr
/usr/include/X11/bitmaps/cntr_ptrmsk
/usr/include/X11/bitmaps/cross_weave
/usr/include/X11/bitmaps/dimple1
/usr/include/X11/bitmaps/dimple3
/usr/include/X11/bitmaps/dot
/usr/include/X11/bitmaps/flagdown
/usr/include/X11/bitmaps/flagup
/usr/include/X11/bitmaps/flipped_gray
/usr/include/X11/bitmaps/gray
/usr/include/X11/bitmaps/gray1
/usr/include/X11/bitmaps/gray3
/usr/include/X11/bitmaps/icon
/usr/include/X11/bitmaps/left_ptr
/usr/include/X11/bitmaps/left_ptrmsk
/usr/include/X11/bitmaps/light_gray
/usr/include/X11/bitmaps/opendot
/usr/include/X11/bitmaps/opendotMask
/usr/include/X11/bitmaps/right_ptr
/usr/include/X11/bitmaps/right_ptrmsk
/usr/include/X11/bitmaps/root_weave
/usr/include/X11/bitmaps/scales
/usr/include/X11/bitmaps/sipb
/usr/include/X11/bitmaps/star
/usr/include/X11/bitmaps/starMask

```

```

/usr/bin/X11/xrefresh
/usr/bin/X11/xseethru
/usr/bin/X11/xset
/usr/bin/X11/xsetroot
/usr/bin/X11/xterm
/usr/bin/X11/xwcreate
/usr/bin/X11/xwd
/usr/bin/X11/xwd2sb
/usr/bin/X11/xwdestroy
/usr/bin/X11/xwininfo
/usr/bin/X11/xwud
/usr/bin/x11start

/usr/include/X11/bitmaps/stipple
/usr/include/X11/bitmaps/target
/usr/include/X11/bitmaps/tie_fighter
/usr/include/X11/bitmaps/wide_weave
/usr/include/X11/bitmaps/wierd_size
/usr/include/X11/bitmaps/wingdogs
/usr/include/X11/bitmaps/woman
/usr/include/X11/bitmaps/xfd_icon
/usr/include/X11/bitmaps/xlogo16
/usr/include/X11/bitmaps/xlogo32
/usr/include/X11/bitmaps/xlogo64
/usr/lib/X11/XErrorDB

```

## Fileset AX11SERV

```

/etc/newconfig/X0devices
/etc/newconfig/X0pointerkeys
/etc/newconfig/X0screens
/etc/newconfig/rgb_dir
/etc/newconfig/rgb.pag
/system/AX11SERV/CDFcustomize
/system/AX11SERV/customize
/system/AX11SERV/revlist
/usr/bin/X11/X
/usr/lib/X11/fonts/12x21apl.scf
/usr/lib/X11/fonts/12x21bas.scf
/usr/lib/X11/fonts/12x28apl.scf
/usr/lib/X11/fonts/12x28bas.scf
/usr/lib/X11/fonts/12x28ibm.scf
/usr/lib/X11/fonts/6x10.scf
/usr/lib/X11/fonts/6x12.scf
/usr/lib/X11/fonts/6x13.scf
/usr/lib/X11/fonts/8x13.scf
/usr/lib/X11/fonts/8x13bold.scf
/usr/lib/X11/fonts/9x15.scf
/usr/lib/X11/fonts/9x16apl.scf
/usr/lib/X11/fonts/9x16bas.scf
/usr/lib/X11/fonts/9x21apl.scf
/usr/lib/X11/fonts/9x21bas.scf
/usr/lib/X11/fonts/9x21ibm.scf
/usr/lib/X11/fonts/a14.scf
/usr/lib/X11/fonts/apl-s25.scf
/usr/lib/X11/fonts/calc.12x16.scf
/usr/lib/X11/fonts/calc.6x8.scf
/usr/lib/X11/fonts/chp-s25.scf

/usr/lib/X11/fonts/grk-s30.scf
/usr/lib/X11/fonts/hbr-s25.scf
/usr/lib/X11/fonts/hbr-s40.scf
/usr/lib/X11/fonts/hp8.10x20.scf
/usr/lib/X11/fonts/hp8.10x20b.scf
/usr/lib/X11/fonts/hp8.12x15.scf
/usr/lib/X11/fonts/hp8.6x13.scf
/usr/lib/X11/fonts/hp8.6x13b.scf
/usr/lib/X11/fonts/hp8.6x8.scf
/usr/lib/X11/fonts/hp8.6x8b.scf
/usr/lib/X11/fonts/hp8.7x10.scf
/usr/lib/X11/fonts/hp8.8x16.scf
/usr/lib/X11/fonts/hp8.8x16b.scf
/usr/lib/X11/fonts/hp8.8x16i.scf
/usr/lib/X11/fonts/ipa-s25.scf
/usr/lib/X11/fonts/iso1.13.scf
/usr/lib/X11/fonts/iso1.13b.scf
/usr/lib/X11/fonts/iso1.15.scf
/usr/lib/X11/fonts/iso1.16.scf
/usr/lib/X11/fonts/iso1.16b.scf
/usr/lib/X11/fonts/iso1.20.scf
/usr/lib/X11/fonts/iso1.20b.scf
/usr/lib/X11/fonts/iso1.8.scf
/usr/lib/X11/fonts/k14.scf
/usr/lib/X11/fonts/kana.10x18.scf
/usr/lib/X11/fonts/kana.10x20.scf
/usr/lib/X11/fonts/kana.12x24.scf
/usr/lib/X11/fonts/kana.8x16.scf
/usr/lib/X11/fonts/kana.8x18.scf
/usr/lib/X11/fonts/kana14.scf

```

```

/usr/lib/X11/fonts/chs-s50.scf      /usr/lib/X11/fonts/krivo.scf
/usr/lib/X11/fonts/cr.12x20.scf    /usr/lib/X11/fonts/lat-s30.scf
/usr/lib/X11/fonts/cr.12x20b.scf   /usr/lib/X11/fonts/line.8x16.scf
/usr/lib/X11/fonts/cursor.scf      /usr/lib/X11/fonts/math.18x30.scf
/usr/lib/X11/fonts/cyr-s25.scf     /usr/lib/X11/fonts/math.6x8.scf
/usr/lib/X11/fonts/cyr-s30.scf     /usr/lib/X11/fonts/math.8x16.scf
/usr/lib/X11/fonts/cyr-s38.scf     /usr/lib/X11/fonts/met25.scf
/usr/lib/X11/fonts/fcor-20.scf     /usr/lib/X11/fonts/micro.scf
/usr/lib/X11/fonts/fg-13.scf       /usr/lib/X11/fonts/oldera.scf
/usr/lib/X11/fonts/fg-16.scf       /usr/lib/X11/fonts/pica.18x30.scf
/usr/lib/X11/fonts/fg-18.scf       /usr/lib/X11/fonts/plunk.scf
/usr/lib/X11/fonts/fg-20.scf       /usr/lib/X11/fonts/r14.scf
/usr/lib/X11/fonts/fg-22.scf       /usr/lib/X11/fonts/rot-s16.scf
/usr/lib/X11/fonts/fg-25.scf       /usr/lib/X11/fonts/sans12.scf
/usr/lib/X11/fonts/fg-30.scf       /usr/lib/X11/fonts/sansb12.scf
/usr/lib/X11/fonts/fg-40.scf       /usr/lib/X11/fonts/sansi12.scf
/usr/lib/X11/fonts/fgi-25.scf      /usr/lib/X11/fonts/serif10.scf
/usr/lib/X11/fonts/fgb-13.scf      /usr/lib/X11/fonts/serif12.scf
/usr/lib/X11/fonts/fgb-25.scf      /usr/lib/X11/fonts/serifb10.scf
/usr/lib/X11/fonts/fgb1-25.scf     /usr/lib/X11/fonts/serifb12.scf
/usr/lib/X11/fonts/fgb1-30.scf     /usr/lib/X11/fonts/serifi10.scf
/usr/lib/X11/fonts/fgi-20.scf      /usr/lib/X11/fonts/serifi12.scf
/usr/lib/X11/fonts/fgi1-25.scf     /usr/lib/X11/fonts/sub.scf
/usr/lib/X11/fonts/fgs-22.scf      /usr/lib/X11/fonts/subsub.scf
/usr/lib/X11/fonts/fixe.scf        /usr/lib/X11/fonts/sup.scf
/usr/lib/X11/fonts/fqxb-25.scf     /usr/lib/X11/fonts/supsup.scf
/usr/lib/X11/fonts/fr-25.scf       /usr/lib/X11/fonts/swd-s30.scf
/usr/lib/X11/fonts/fr-33.scf       /usr/lib/X11/fonts/sym-s25.scf
/usr/lib/X11/fonts/fr1-25.scf      /usr/lib/X11/fonts/sym-s53.scf
/usr/lib/X11/fonts/fr2-25.scf      /usr/lib/X11/fonts/variable.scf
/usr/lib/X11/fonts/fr3-25.scf      /usr/lib/X11/fonts/vbee-36.scf
/usr/lib/X11/fonts/frb-32.scf      /usr/lib/X11/fonts/vct1-25.scf
/usr/lib/X11/fonts/fri-33.scf      /usr/lib/X11/fonts/vg-13.scf
/usr/lib/X11/fonts/fri1-25.scf     /usr/lib/X11/fonts/vg-20.scf
/usr/lib/X11/fonts/ger-s35.scf     /usr/lib/X11/fonts/vg-25.scf
/usr/lib/X11/fonts/grk-s25.scf     /usr/lib/X11/fonts/vg-31.scf
/usr/lib/X11/fonts/vg-40.scf       /usr/lib/X11/fonts/vtbold.scf
/usr/lib/X11/fonts/vgb-25.scf      /usr/lib/X11/fonts/vtsingle.scf
/usr/lib/X11/fonts/vgb-31.scf      /usr/lib/X11/fonts/vxms-37.scf
/usr/lib/X11/fonts/vgbc-25.scf     /usr/lib/X11/fonts/vxms-43.scf
/usr/lib/X11/fonts/vgh-25.scf      /usr/lib/X11/fonts/xif-s25.scf
/usr/lib/X11/fonts/vgi-20.scf      /usr/lib/X11/fonts/vgi-25.scf
/usr/lib/X11/fonts/vgi-31.scf      /usr/lib/X11/fonts/vgl-40.scf
/usr/lib/X11/fonts/vgvb-31.scf     /usr/lib/X11/fonts/vmic-25.scf
/usr/lib/X11/fonts/vr-20.scf       /usr/lib/X11/fonts/vr-25.scf

```

/usr/lib/X11/fonts/vr-27.scf	/usr/lib/X11/fonts/vr-30.scf
/usr/lib/X11/fonts/vr-31.scf	/usr/lib/X11/fonts/vr-40.scf
/usr/lib/X11/fonts/vrb-25.scf	/usr/lib/X11/fonts/vrb-30.scf
/usr/lib/X11/fonts/vrb-31.scf	/usr/lib/X11/fonts/vrb-35.scf
/usr/lib/X11/fonts/vrb-37.scf	/usr/lib/X11/fonts/vri-25.scf
/usr/lib/X11/fonts/vri-30.scf	/usr/lib/X11/fonts/vri-31.scf
/usr/lib/X11/fonts/vri-40.scf	/usr/lib/X11/fonts/vsg-114.scf
/usr/lib/X11/fonts/vsgn-57.scf	/usr/lib/X11/fonts/vshd-40.scf
/usr/lib/nls/katakana/X11font/alt.high.scf	
/usr/lib/nls/katakana/X11font/alt.high.snf	
/usr/lib/nls/katakana/X11font/alt.low.scf	
/usr/lib/nls/katakana/X11font/alt.low.snf	
/usr/lib/nls/katakana/X11font/alt.med.scf	
/usr/lib/nls/katakana/X11font/alt.med.snf	
/usr/lib/nls/katakana/X11font/base.high.scf	
/usr/lib/nls/katakana/X11font/base.high.snf	
/usr/lib/nls/katakana/X11font/base.low.scf	
/usr/lib/nls/katakana/X11font/base.low.snf	
/usr/lib/nls/katakana/X11font/base.med.scf	
/usr/lib/nls/katakana/X11font/base.med.snf	
/usr/lib/nls/n-computer/X11font/alt.high.scf	
/usr/lib/nls/n-computer/X11font/alt.high.snf	
/usr/lib/nls/n-computer/X11font/alt.low.scf	
/usr/lib/nls/n-computer/X11font/alt.low.snf	
/usr/lib/nls/n-computer/X11font/alt.med.scf	
/usr/lib/nls/n-computer/X11font/alt.med.snf	
/usr/lib/nls/n-computer/X11font/base.high.scf	
/usr/lib/nls/n-computer/X11font/base.high.snf	
/usr/lib/nls/n-computer/X11font/base.low.scf	
/usr/lib/nls/n-computer/X11font/base.low.snf	
/usr/lib/nls/n-computer/X11font/base.med.scf	
/usr/lib/nls/n-computer/X11font/base.med.snf	

## Partition WINDOW

### Fileset AWINDOW

/system/AWINDOW/CDFcustomize	/usr/bin/wdisp
/system/AWINDOW/customize	/usr/bin/wfont
/system/AWINDOW/revlist	/usr/bin/wlist
/usr/bin/wborder	/usr/bin/wmove
/usr/bin/wmstart	/usr/bin/wmready
/usr/lib/gserver	/usr/bin/wmstop
/usr/lib/stserver	/usr/bin/wselect

/usr/lib/wm  
usr/bin/wcreate  
usr/bin/wdestroy

/usr/bin/wsh  
/usr/bin/wsize  
/usr/lib/t0server

## Partition PROG\_LANGS

### Fileset AC

/bin/cc  
/lib/c210  
/lib/ccom10  
/lib/cpp

/lib/libc.a  
/system/AC/revlist  
/usr/lib/libV3.a  
/usr/lib/libsetjmp.a

### Fileset AFPFA

/etc/fpocode  
/etc/fpload

/system/AFPFA/customize  
/system/AFPFA/revlist

### Fileset APROG

/bin/ar  
/bin/as  
/bin/as10  
/bin/ld

/bin/make  
/bin/strip  
/system/APROG/revlist  
/usr/lib/ar.ranlib

## Partition MISC\_UTILS

### Fileset AAUDIT

/etc/newconfig/auditrc  
/etc/tsconvert  
/system/AAUDIT/AAAUDIT  
/system/AAUDIT/revlist  
/usr/bin/audevent

/usr/bin/audisp  
/usr/bin/audomon  
/usr/bin/audsys  
/usr/bin/audusr

### Fileset ACONFIG

/etc/conf/dfile.cnode  
/etc/conf/dfile.cnodemax  
/etc/conf/dfile.cnodemin  
/etc/conf/dfile.full  
/etc/conf/dfile.full.lan  
/etc/conf/dfile.maxservr

/etc/conf/h/types.h  
/etc/conf/h/user.h  
/etc/conf/h/vnode.h  
/etc/conf/h/x25\_diag.h  
/etc/conf/libdevelop.a  
/etc/conf/libdil\_srm.a

/etc/conf/dfile.min	/etc/conf/libdreq.a
/etc/conf/dfile.minserverv	/etc/conf/libdskless.a
/etc/conf/h/acl.h	/etc/conf/libkreq.a
/etc/conf/h/buf.h	/etc/conf/libmin.a
/etc/conf/h/callout.h	/etc/conf/libserver.a
/etc/conf/h/cmap.h	/etc/conf/libsysV.a
/etc/conf/h/conf.h	/etc/conf/libvm.a
/etc/conf/h/devices.h	/etc/conf/machine/a.out.h
/etc/conf/h/dir.h	/etc/conf/machine/dilio.h
/etc/conf/h/dmap.h	/etc/conf/machine/dm.h
/etc/conf/h/dnlc.h	/etc/conf/machine/drv_lan0.h
/etc/conf/h/dux_mbuf.h	/etc/conf/machine/drvhw.h
/etc/conf/h/errno.h	/etc/conf/machine/dux_hooks.h
/etc/conf/h/fsdir.h	/etc/conf/machine/duxparam.h
/etc/conf/h/ino.h	/etc/conf/machine/hpibio.h
/etc/conf/h/inode.h	/etc/conf/machine/if.h
/etc/conf/h/ioctl.h	/etc/conf/machine/if_ether.h
/etc/conf/h/ipc.h	/etc/conf/machine/in.h
/etc/conf/h/kernel.h	/etc/conf/machine/intrpt.h
/etc/conf/h/kpreempt.h	/etc/conf/machine/iobuf.h
/etc/conf/h/magic.h	/etc/conf/machine/lnatypes.h
/etc/conf/h/map.h	/etc/conf/machine/nsdiag0.h
/etc/conf/h/mbuf.h	/etc/conf/machine/nsp.h
/etc/conf/h/msg.h	/etc/conf/machine/param.h
/etc/conf/h/ns_diag.h	/etc/conf/machine/pcb.h
/etc/conf/h/opt.h	/etc/conf/machine/protocol.h
/etc/conf/h/param.h	/etc/conf/machine/pte.h
/etc/conf/h/privgrp.h	/etc/conf/machine/rmswap.h
/etc/conf/h/proc.h	/etc/conf/machine/sitemap.h
/etc/conf/h/pty.h	/etc/conf/machine/timeout.h
/etc/conf/h/ptyio.h	/etc/config
/etc/conf/h/resource.h	/etc/mkrs
/etc/conf/h/rfa_user.h	/etc/mkrs.devs
/etc/conf/h/sem.h	/etc/mkrs.init
/etc/conf/h/shm.h	/etc/mkrs.mknod
/etc/conf/h/signal.h	/etc/mkrs.mount
/etc/conf/h/socket.h	/etc/mkrs.swap
/etc/conf/h/space.h	/etc/mkrs.tool
/etc/conf/h/sysmacros.h	/etc/mkrs.umount
/etc/conf/h/systm.h	/etc/newconfig/master
/etc/conf/h/text.h	/system/ACONFIG/CDFcustomize
/etc/conf/h/time.h	/system/ACONFIG/customize
/etc/conf/h/tty.h	/system/ACONFIG/revlist

## Fileset ATRMINFO

/system/ATRMINFO/ATRMINFO.links	/usr/lib/terminfo/4/4424-el-2
/system/ATRMINFO/customize	/usr/lib/terminfo/4/45
/system/ATRMINFO/revlist	/usr/lib/terminfo/4/450
/usr/lib/terminfo/1/100	/usr/lib/terminfo/5/5520
/usr/lib/terminfo/1/110plus	/usr/lib/terminfo/7/70092
/usr/lib/terminfo/1/1520	/usr/lib/terminfo/7/70092-w
/usr/lib/terminfo/1/1521	/usr/lib/terminfo/7/70092A
/usr/lib/terminfo/1/1620	/usr/lib/terminfo/7/70092A-w
/usr/lib/terminfo/1/1620-m8	/usr/lib/terminfo/7/70092a
/usr/lib/terminfo/1/1640	/usr/lib/terminfo/7/70092a-w
/usr/lib/terminfo/1/1640-m8	/usr/lib/terminfo/7/70094
/usr/lib/terminfo/1/1700	/usr/lib/terminfo/7/70094-w
/usr/lib/terminfo/1/11line	/usr/lib/terminfo/7/70094A
/usr/lib/terminfo/1/11linepty	/usr/lib/terminfo/7/70094A-w
/usr/lib/terminfo/2/2500	/usr/lib/terminfo/7/70094a
/usr/lib/terminfo/2/2621P	/usr/lib/terminfo/7/70094a-w
/usr/lib/terminfo/2/2621p	/usr/lib/terminfo/7/735
/usr/lib/terminfo/3/30	/usr/lib/terminfo/7/743
/usr/lib/terminfo/3/300	/usr/lib/terminfo/7/745
/usr/lib/terminfo/3/300s	/usr/lib/terminfo/8/8001
/usr/lib/terminfo/3/3045	/usr/lib/terminfo/8/8510
/usr/lib/terminfo/3/31	/usr/lib/terminfo/9/905
/usr/lib/terminfo/3/3101	/usr/lib/terminfo/9/910+
/usr/lib/terminfo/3/33	/usr/lib/terminfo/9/912
/usr/lib/terminfo/3/333	/usr/lib/terminfo/9/912-2p
/usr/lib/terminfo/3/35	/usr/lib/terminfo/9/9122p
/usr/lib/terminfo/3/35714	/usr/lib/terminfo/9/912b
/usr/lib/terminfo/3/35714a	/usr/lib/terminfo/9/912c
/usr/lib/terminfo/3/36	/usr/lib/terminfo/9/912cc
/usr/lib/terminfo/3/37	/usr/lib/terminfo/9/920
/usr/lib/terminfo/3/380	/usr/lib/terminfo/9/920-2p
/usr/lib/terminfo/3/382	/usr/lib/terminfo/9/9202p
/usr/lib/terminfo/3/3a	/usr/lib/terminfo/9/920b
/usr/lib/terminfo/3/3a+	/usr/lib/terminfo/9/920c
/usr/lib/terminfo/4/40	/usr/lib/terminfo/9/925
/usr/lib/terminfo/4/400	/usr/lib/terminfo/9/925E
/usr/lib/terminfo/4/4012	/usr/lib/terminfo/9/950
/usr/lib/terminfo/4/4013	/usr/lib/terminfo/9/950-2p
/usr/lib/terminfo/4/4014	/usr/lib/terminfo/9/950-4p
/usr/lib/terminfo/4/4014-sm	/usr/lib/terminfo/9/950-rv
/usr/lib/terminfo/4/4015	/usr/lib/terminfo/9/950-rv-2p
/usr/lib/terminfo/4/4015-sm	/usr/lib/terminfo/9/950-rv-4p
/usr/lib/terminfo/4/4023	/usr/lib/terminfo/9/9502p



/usr/lib/terminfo/4/4024	/usr/lib/terminfo/9/9504p
/usr/lib/terminfo/4/4025	/usr/lib/terminfo/9/950rv
/usr/lib/terminfo/4/4025-17	/usr/lib/terminfo/9/950rv2p
/usr/lib/terminfo/4/4025-17ws	/usr/lib/terminfo/9/950rv4p
/usr/lib/terminfo/4/4025cu	/usr/lib/terminfo/9/9835
/usr/lib/terminfo/4/4025ex	/usr/lib/terminfo/9/9845
/usr/lib/terminfo/4/4027	/usr/lib/terminfo/D/D1182A
/usr/lib/terminfo/4/4027-17	/usr/lib/terminfo/D/D1182a
/usr/lib/terminfo/4/4027-17ws	/usr/lib/terminfo/a/a980
/usr/lib/terminfo/4/4027cu	/usr/lib/terminfo/a/aa
/usr/lib/terminfo/4/4027ex	/usr/lib/terminfo/a/aaa
/usr/lib/terminfo/4/4080	/usr/lib/terminfo/a/aaa-18
/usr/lib/terminfo/4/4112	/usr/lib/terminfo/a/aaa-18-rv
/usr/lib/terminfo/4/4112-5	/usr/lib/terminfo/a/aaa-20
/usr/lib/terminfo/4/4112-d	/usr/lib/terminfo/a/aaa-22
/usr/lib/terminfo/4/4112-nd	/usr/lib/terminfo/a/aaa-24
/usr/lib/terminfo/4/4113	/usr/lib/terminfo/a/aaa-24-rv
/usr/lib/terminfo/4/4114	/usr/lib/terminfo/a/aaa-26
/usr/lib/terminfo/4/42	/usr/lib/terminfo/a/aaa-28
/usr/lib/terminfo/4/42-nl	/usr/lib/terminfo/a/aaa-29
/usr/lib/terminfo/4/43	/usr/lib/terminfo/a/aaa-29-ctxt
/usr/lib/terminfo/4/4424	/usr/lib/terminfo/a/aaa-29-np
/usr/lib/terminfo/4/4424-2	/usr/lib/terminfo/a/aaa-29-rv
/usr/lib/terminfo/a/aaa-30	/usr/lib/terminfo/a/alto
/usr/lib/terminfo/a/aaa-30-ctxt	/usr/lib/terminfo/a/altoh19
/usr/lib/terminfo/a/aaa-30-rv	/usr/lib/terminfo/a/altoheath
/usr/lib/terminfo/a/aaa-30-rv-ctxt	/usr/lib/terminfo/a/ambas
/usr/lib/terminfo/a/aaa-30-s	/usr/lib/terminfo/a/ambassador
/usr/lib/terminfo/a/aaa-30-s-ctxt	/usr/lib/terminfo/a/ampex
/usr/lib/terminfo/a/aaa-30-s-rv	/usr/lib/terminfo/a/annarbor
/usr/lib/terminfo/a/aaa-30-s-rv-ct	/usr/lib/terminfo/a/apple
/usr/lib/terminfo/a/aaa-36	/usr/lib/terminfo/a/arpanet
/usr/lib/terminfo/a/aaa-36-rv	/usr/lib/terminfo/b/bantam
/usr/lib/terminfo/a/aaa-40	/usr/lib/terminfo/b/bc
/usr/lib/terminfo/a/aaa-40-rv	/usr/lib/terminfo/b/beacon
/usr/lib/terminfo/a/aaa-48	/usr/lib/terminfo/b/beehiveIIIm
/usr/lib/terminfo/a/aaa-48-rv	/usr/lib/terminfo/b/bg
/usr/lib/terminfo/a/aaa-59	/usr/lib/terminfo/b/bg-ni
/usr/lib/terminfo/a/aaa-60	/usr/lib/terminfo/b/bg-nv
/usr/lib/terminfo/a/aaa-60-rv	/usr/lib/terminfo/b/bg-rv
/usr/lib/terminfo/a/aaa-60-s	/usr/lib/terminfo/b/bg1.25
/usr/lib/terminfo/a/aaa-60-s-rv	/usr/lib/terminfo/b/bg1.25-nv

/usr/lib/terminfo/a/aaa-ctxt	/usr/lib/terminfo/b/bg1.25-rv
/usr/lib/terminfo/a/aaa-db	/usr/lib/terminfo/b/bg2.0
/usr/lib/terminfo/a/aaa-rv	/usr/lib/terminfo/b/bg2.0-nv
/usr/lib/terminfo/a/aaa-rv-ctxt	/usr/lib/terminfo/b/bg2.0-rv
/usr/lib/terminfo/a/aaa-rv-unk	/usr/lib/terminfo/b/bh3m
/usr/lib/terminfo/a/aaa-s	/usr/lib/terminfo/b/bitgraph
/usr/lib/terminfo/a/aaa-s-rv	/usr/lib/terminfo/b/bitgraph-ni
/usr/lib/terminfo/a/aaa-unk	/usr/lib/terminfo/b/bitgraph-nv
/usr/lib/terminfo/a/aaa18	/usr/lib/terminfo/b/bitgraph-rv
/usr/lib/terminfo/a/aaa20	/usr/lib/terminfo/b/blit
/usr/lib/terminfo/a/aaa22	/usr/lib/terminfo/b/blit-pb
/usr/lib/terminfo/a/aaa24	/usr/lib/terminfo/b/blitlayer
/usr/lib/terminfo/a/aaa26	/usr/lib/terminfo/b/bussiplexer
/usr/lib/terminfo/a/aaa28	/usr/lib/terminfo/c/c100
/usr/lib/terminfo/a/aaa29	/usr/lib/terminfo/c/c100-1p
/usr/lib/terminfo/a/aaa30	/usr/lib/terminfo/c/c100-4p
/usr/lib/terminfo/a/aaa36	/usr/lib/terminfo/c/c100-rv
/usr/lib/terminfo/a/aaa40	/usr/lib/terminfo/c/c100-rv-4p
/usr/lib/terminfo/a/aaa48	/usr/lib/terminfo/c/c100-rv-4p-na
/usr/lib/terminfo/a/aaa59	/usr/lib/terminfo/c/c100-rv-4p-pp
/usr/lib/terminfo/a/aaa60	/usr/lib/terminfo/c/c100-rv-na
/usr/lib/terminfo/a/aaadb	/usr/lib/terminfo/c/c100-rv-pp
/usr/lib/terminfo/a/act4	/usr/lib/terminfo/c/c1004p
/usr/lib/terminfo/a/act5	/usr/lib/terminfo/c/c100rv
/usr/lib/terminfo/a/act5s	/usr/lib/terminfo/c/c100rv4p
/usr/lib/terminfo/a/ad	/usr/lib/terminfo/c/c100rv4pna
/usr/lib/terminfo/a/adrrinfo	/usr/lib/terminfo/c/c100rv4ppp
/usr/lib/terminfo/a/adm1	/usr/lib/terminfo/c/c100rvna
/usr/lib/terminfo/a/adm1a	/usr/lib/terminfo/c/c100rvpp
/usr/lib/terminfo/a/adm2	/usr/lib/terminfo/c/c100rvs
/usr/lib/terminfo/a/adm21	/usr/lib/terminfo/c/c100s
/usr/lib/terminfo/a/adm3	/usr/lib/terminfo/c/c104
/usr/lib/terminfo/a/adm31	/usr/lib/terminfo/c/c108
/usr/lib/terminfo/a/adm36	/usr/lib/terminfo/c/c108-4
/usr/lib/terminfo/a/adm3a	/usr/lib/terminfo/c/c108-4p
/usr/lib/terminfo/a/adm3a+	/usr/lib/terminfo/c/c108-8
/usr/lib/terminfo/a/adm3aplus	/usr/lib/terminfo/c/c108-8p
/usr/lib/terminfo/a/adm42	/usr/lib/terminfo/c/c108-8p-na
/usr/lib/terminfo/a/adm42-nl	/usr/lib/terminfo/c/c108-8p-rv-na
/usr/lib/terminfo/a/adm5	/usr/lib/terminfo/c/c108-na
/usr/lib/terminfo/a/aed	/usr/lib/terminfo/c/c108-na-8p
/usr/lib/terminfo/a/aed512	/usr/lib/terminfo/c/c108-rv

/usr/lib/terminfo/a/agile	/usr/lib/terminfo/c/c108-rv-4p
/usr/lib/terminfo/a/agiles	/usr/lib/terminfo/c/c108-rv-8p
/usr/lib/terminfo/a/aj	/usr/lib/terminfo/c/c108-rv-na
/usr/lib/terminfo/a/aj830	/usr/lib/terminfo/c/c108-rv-na-8p
/usr/lib/terminfo/a/aj832	/usr/lib/terminfo/c/c108-w
/usr/lib/terminfo/c/c108-w-8p	/usr/lib/terminfo/d/dmdt80-w
/usr/lib/terminfo/c/ca	/usr/lib/terminfo/d/dp3
/usr/lib/terminfo/c/ca22851	/usr/lib/terminfo/d/dp3360
/usr/lib/terminfo/c/carlock	/usr/lib/terminfo/d/ds40
/usr/lib/terminfo/c/cb-unix	/usr/lib/terminfo/d/ds40-2
/usr/lib/terminfo/c/cbblit	/usr/lib/terminfo/d/dt80
/usr/lib/terminfo/c/cbunix	/usr/lib/terminfo/d/dt80-w
/usr/lib/terminfo/c/cdc	/usr/lib/terminfo/d/dtc
/usr/lib/terminfo/c/cdc456	/usr/lib/terminfo/d/dtc300s
/usr/lib/terminfo/c/cdc456tst	/usr/lib/terminfo/d/dtc382
/usr/lib/terminfo/c/cdi	/usr/lib/terminfo/d/dw
/usr/lib/terminfo/c/cdi1203	/usr/lib/terminfo/d/dw1
/usr/lib/terminfo/c/cg7900	/usr/lib/terminfo/d/dw2
/usr/lib/terminfo/c/chromaticcs	/usr/lib/terminfo/d/dw3
/usr/lib/terminfo/c/ci8510	/usr/lib/terminfo/d/dw4
/usr/lib/terminfo/c/citoh	/usr/lib/terminfo/e/env230
/usr/lib/terminfo/c/citoh-6lpi	/usr/lib/terminfo/e/envision230
/usr/lib/terminfo/c/citoh-8lpi	/usr/lib/terminfo/e/ep40
/usr/lib/terminfo/c/citoh-comp	/usr/lib/terminfo/e/ep4000
/usr/lib/terminfo/c/citoh-elite	/usr/lib/terminfo/e/ep4080
/usr/lib/terminfo/c/citoh-pica	/usr/lib/terminfo/e/ep48
/usr/lib/terminfo/c/citoh-prop	/usr/lib/terminfo/e/esprit
/usr/lib/terminfo/c/citoh-ps	/usr/lib/terminfo/e/ethernet
/usr/lib/terminfo/c/compucolor2	/usr/lib/terminfo/e/ex3000
/usr/lib/terminfo/c/compucolorii	/usr/lib/terminfo/e/exidy
/usr/lib/terminfo/c/concept	/usr/lib/terminfo/e/exidy2500
/usr/lib/terminfo/c/concept100	/usr/lib/terminfo/f/f100
/usr/lib/terminfo/c/concept100-rv	/usr/lib/terminfo/f/f100-rv
/usr/lib/terminfo/c/concept100-rv-	/usr/lib/terminfo/f/f1720
/usr/lib/terminfo/c/concept108-4p	/usr/lib/terminfo/f/f1720a
/usr/lib/terminfo/c/concept108-8p	/usr/lib/terminfo/f/falco
/usr/lib/terminfo/c/concept108-na-	/usr/lib/terminfo/f/falco-p
/usr/lib/terminfo/c/concept108-rv-	/usr/lib/terminfo/f/fixterm
/usr/lib/terminfo/c/concept108-w-8	/usr/lib/terminfo/f/fox
/usr/lib/terminfo/c/ct82	/usr/lib/terminfo/f/freedom
/usr/lib/terminfo/c/ct8500	/usr/lib/terminfo/f/freedom-rv
/usr/lib/terminfo/d/d100	/usr/lib/terminfo/f/freedom100
/usr/lib/terminfo/d/d132	/usr/lib/terminfo/g/gigi

/usr/lib/terminfo/d/d200	/usr/lib/terminfo/g/gsi
/usr/lib/terminfo/d/d80	/usr/lib/terminfo/g/gt100
/usr/lib/terminfo/d/d800	/usr/lib/terminfo/g/gt100a
/usr/lib/terminfo/d/datagraphix	/usr/lib/terminfo/g/gt40
/usr/lib/terminfo/d/datamedia2500	/usr/lib/terminfo/g/gt42
/usr/lib/terminfo/d/datapoint	/usr/lib/terminfo/h/h1000
/usr/lib/terminfo/d/dataspeed40	/usr/lib/terminfo/h/h1420
/usr/lib/terminfo/d/dd5000	/usr/lib/terminfo/h/h1500
/usr/lib/terminfo/d/debug	/usr/lib/terminfo/h/h1510
/usr/lib/terminfo/d/decwriter	/usr/lib/terminfo/h/h1520
/usr/lib/terminfo/d/delta	/usr/lib/terminfo/h/h1552
/usr/lib/terminfo/d/dg	/usr/lib/terminfo/h/h1552-rv
/usr/lib/terminfo/d/dg6053	/usr/lib/terminfo/h/h19
/usr/lib/terminfo/d/diablo	/usr/lib/terminfo/h/h19-a
/usr/lib/terminfo/d/dialogue	/usr/lib/terminfo/h/h19-b
/usr/lib/terminfo/d/dialogue80	/usr/lib/terminfo/h/h19-bs
/usr/lib/terminfo/d/dialup	/usr/lib/terminfo/h/h19-pb
/usr/lib/terminfo/d/digilog	/usr/lib/terminfo/h/h19-smul
/usr/lib/terminfo/d/direct	/usr/lib/terminfo/h/h19-u
/usr/lib/terminfo/d/direct800	/usr/lib/terminfo/h/h19A
/usr/lib/terminfo/d/dm1520	/usr/lib/terminfo/h/h19a
/usr/lib/terminfo/d/dm1521	/usr/lib/terminfo/h/h19b
/usr/lib/terminfo/d/dm2500	/usr/lib/terminfo/h/h19bs
/usr/lib/terminfo/d/dm3025	/usr/lib/terminfo/h/h19u
/usr/lib/terminfo/d/dm3045	/usr/lib/terminfo/h/h2000
/usr/lib/terminfo/d/dm80	/usr/lib/terminfo/h/h89
/usr/lib/terminfo/d/dm80-w	/usr/lib/terminfo/h/heath
/usr/lib/terminfo/d/dmdt80	/usr/lib/terminfo/h/heath-19
/usr/lib/terminfo/h/heath-ansi	/usr/lib/terminfo/m/mw2
/usr/lib/terminfo/h/heathkit	/usr/lib/terminfo/n/nec
/usr/lib/terminfo/h/heathkit-a	/usr/lib/terminfo/n/netronics
/usr/lib/terminfo/h/hk	/usr/lib/terminfo/n/network
/usr/lib/terminfo/h/hp2621p	/usr/lib/terminfo/n/netx
/usr/lib/terminfo/h/hp35714	/usr/lib/terminfo/n/nuc
/usr/lib/terminfo/h/hp35714a	/usr/lib/terminfo/n/nucterm
/usr/lib/terminfo/h/hp9835	/usr/lib/terminfo/o/o31
/usr/lib/terminfo/h/hp9845	/usr/lib/terminfo/o/oadm31
/usr/lib/terminfo/h/hpD1182	/usr/lib/terminfo/o/obitgraph
/usr/lib/terminfo/h/hpD1182A	/usr/lib/terminfo/o/obitgraph-nv
/usr/lib/terminfo/h/hpD1182a	/usr/lib/terminfo/o/obitgraph-rv
/usr/lib/terminfo/i/i100	/usr/lib/terminfo/o/oblit
/usr/lib/terminfo/i/i3101	/usr/lib/terminfo/o/oc100
/usr/lib/terminfo/i/i400	/usr/lib/terminfo/o/oconcept

/usr/lib/terminfo/i/ibm	/usr/lib/terminfo/o/ojerq
/usr/lib/terminfo/i/ibm3101	/usr/lib/terminfo/o/omron
/usr/lib/terminfo/i/infoton	/usr/lib/terminfo/o/ovi300
/usr/lib/terminfo/i/intertec	/usr/lib/terminfo/o/owl
/usr/lib/terminfo/i/intertube	/usr/lib/terminfo/p/patch
/usr/lib/terminfo/i/intertube2	/usr/lib/terminfo/p/patchboard
/usr/lib/terminfo/i/intext	/usr/lib/terminfo/p/pbox
/usr/lib/terminfo/i/intext2	/usr/lib/terminfo/p/pc
/usr/lib/terminfo/i/intextii	/usr/lib/terminfo/p/pe550
/usr/lib/terminfo/i/ips	/usr/lib/terminfo/p/plasma
/usr/lib/terminfo/i/iq120	/usr/lib/terminfo/p/plugboard
/usr/lib/terminfo/i/iq140	/usr/lib/terminfo/p/plus
/usr/lib/terminfo/i/isc	/usr/lib/terminfo/p/print
/usr/lib/terminfo/i/isc8001	/usr/lib/terminfo/p/printer
/usr/lib/terminfo/i/it	/usr/lib/terminfo/p/printerbox
/usr/lib/terminfo/i/it2	/usr/lib/terminfo/p/printing
/usr/lib/terminfo/j/jjerq	/usr/lib/terminfo/p/ps
/usr/lib/terminfo/k/kclc	/usr/lib/terminfo/p/pty
/usr/lib/terminfo/k/kctm	/usr/lib/terminfo/q/qume
/usr/lib/terminfo/l/la120	/usr/lib/terminfo/q/qume5
/usr/lib/terminfo/l/layer	/usr/lib/terminfo/q/qvt101
/usr/lib/terminfo/l/lp	/usr/lib/terminfo/r/ramtek
/usr/lib/terminfo/l/lpr	/usr/lib/terminfo/r/rayterm
/usr/lib/terminfo/m/mdl110	/usr/lib/terminfo/r/reach
/usr/lib/terminfo/m/megatek	/usr/lib/terminfo/r/regent
/usr/lib/terminfo/m/microb	/usr/lib/terminfo/r/regent100
/usr/lib/terminfo/m/microbee	/usr/lib/terminfo/r/regent20
/usr/lib/terminfo/m/microkit	/usr/lib/terminfo/r/regent200
/usr/lib/terminfo/m/microterm	/usr/lib/terminfo/r/regent25
/usr/lib/terminfo/m/microterm5	/usr/lib/terminfo/r/regent40
/usr/lib/terminfo/m/mime	/usr/lib/terminfo/r/regent40+
/usr/lib/terminfo/m/mime-3a	/usr/lib/terminfo/r/regent40-s
/usr/lib/terminfo/m/mime-3ax	/usr/lib/terminfo/r/regent60
/usr/lib/terminfo/m/mime-adm3a	/usr/lib/terminfo/r/regent60-na
/usr/lib/terminfo/m/mime-adm3ax	/usr/lib/terminfo/s/s1500
/usr/lib/terminfo/m/mime-fb	/usr/lib/terminfo/s/sb1
/usr/lib/terminfo/m/mime-hb	/usr/lib/terminfo/s/sb2
/usr/lib/terminfo/m/mime1	/usr/lib/terminfo/s/sb3
/usr/lib/terminfo/m/mime2	/usr/lib/terminfo/s/sbg
/usr/lib/terminfo/m/mime2a	/usr/lib/terminfo/s/sbi
/usr/lib/terminfo/m/mime2a-s	/usr/lib/terminfo/s/screwpoint
/usr/lib/terminfo/m/mime2a-v	/usr/lib/terminfo/s/sexidy
/usr/lib/terminfo/m/mime2as	/usr/lib/terminfo/s/smarterm

```

/usr/lib/terminfo/m/mime2av
/usr/lib/terminfo/m/mimefb
/usr/lib/terminfo/m/mimehb
/usr/lib/terminfo/m/mimeib
/usr/lib/terminfo/m/mimeii
/usr/lib/terminfo/m/minansi
/usr/lib/terminfo/m/mkt
/usr/lib/terminfo/m/mransi
/usr/lib/terminfo/s/sun1
/usr/lib/terminfo/s/superbee
/usr/lib/terminfo/s/superbrain
/usr/lib/terminfo/s/switch
/usr/lib/terminfo/s/swtp
/usr/lib/terminfo/s/synertek
/usr/lib/terminfo/s/system1
/usr/lib/terminfo/t/t10
/usr/lib/terminfo/t/t1061
/usr/lib/terminfo/t/t1061f
/usr/lib/terminfo/t/t16
/usr/lib/terminfo/t/t3700
/usr/lib/terminfo/t/t3800
/usr/lib/terminfo/t/t500
/usr/lib/terminfo/t/tab
/usr/lib/terminfo/t/tab132
/usr/lib/terminfo/t/tab132-rv
/usr/lib/terminfo/t/tab132-w
/usr/lib/terminfo/t/tab132-w-rv
/usr/lib/terminfo/t/tec
/usr/lib/terminfo/t/tec400
/usr/lib/terminfo/t/tec500
/usr/lib/terminfo/t/tek
/usr/lib/terminfo/t/tek4012
/usr/lib/terminfo/t/tek4013
/usr/lib/terminfo/t/tek4014
/usr/lib/terminfo/t/tek4014-sm
/usr/lib/terminfo/t/tek4015
/usr/lib/terminfo/t/tek4015-sm
/usr/lib/terminfo/t/tek4023
/usr/lib/terminfo/t/tek4024
/usr/lib/terminfo/t/tek4025
/usr/lib/terminfo/t/tek4027
/usr/lib/terminfo/t/tek4112
/usr/lib/terminfo/t/teleray
/usr/lib/terminfo/s/smarterm-s
/usr/lib/terminfo/s/smartvid
/usr/lib/terminfo/s/sol
/usr/lib/terminfo/s/sol1
/usr/lib/terminfo/s/sol2
/usr/lib/terminfo/s/soroc
/usr/lib/terminfo/s/spinwriter
/usr/lib/terminfo/s/sun
/usr/lib/terminfo/t/tvi905
/usr/lib/terminfo/t/tvi910+
/usr/lib/terminfo/t/tvi912
/usr/lib/terminfo/t/tvi912-2p
/usr/lib/terminfo/t/tvi9122p
/usr/lib/terminfo/t/tvi912b
/usr/lib/terminfo/t/tvi912c
/usr/lib/terminfo/t/tvi912cc
/usr/lib/terminfo/t/tvi912cup@
/usr/lib/terminfo/t/tvi920
/usr/lib/terminfo/t/tvi920-2p
/usr/lib/terminfo/t/tvi9202p
/usr/lib/terminfo/t/tvi920b
/usr/lib/terminfo/t/tvi920c
/usr/lib/terminfo/t/tvi925
/usr/lib/terminfo/t/tvi925E
/usr/lib/terminfo/t/tvi950
/usr/lib/terminfo/t/tvi950-2p
/usr/lib/terminfo/t/tvi950-4p
/usr/lib/terminfo/t/tvi950-ap
/usr/lib/terminfo/t/tvi950-b
/usr/lib/terminfo/t/tvi950-ns
/usr/lib/terminfo/t/tvi950-rv
/usr/lib/terminfo/t/tvi950-rv-2p
/usr/lib/terminfo/t/tvi950-rv-4p
/usr/lib/terminfo/t/tvi9502p
/usr/lib/terminfo/t/tvi9504p
/usr/lib/terminfo/t/tvi950b
/usr/lib/terminfo/t/tvi950ns
/usr/lib/terminfo/t/tvi950rv
/usr/lib/terminfo/t/tvi950rv2p
/usr/lib/terminfo/t/tvi950rv4p
/usr/lib/terminfo/u/ubell
/usr/lib/terminfo/u/ubellchar
/usr/lib/terminfo/u/unitrm18

```

/usr/lib/terminfo/t/teletec	/usr/lib/terminfo/v/vc103
/usr/lib/terminfo/t/televideo950	/usr/lib/terminfo/v/vc203
/usr/lib/terminfo/t/terak	/usr/lib/terminfo/v/vc303
/usr/lib/terminfo/t/terminet	/usr/lib/terminfo/v/vc303-a
/usr/lib/terminfo/t/terminet1200	/usr/lib/terminfo/v/vc403a
/usr/lib/terminfo/t/terminet300	/usr/lib/terminfo/v/vc404
/usr/lib/terminfo/t/tex	/usr/lib/terminfo/v/vc404-na
/usr/lib/terminfo/t/ti	/usr/lib/terminfo/v/vc404-s
/usr/lib/terminfo/t/ti700	/usr/lib/terminfo/v/vc404-s-na
/usr/lib/terminfo/t/ti733	/usr/lib/terminfo/v/vc415
/usr/lib/terminfo/t/ti735	/usr/lib/terminfo/v/vi200
/usr/lib/terminfo/t/ti745	/usr/lib/terminfo/v/vi200-f
/usr/lib/terminfo/t/ti800	/usr/lib/terminfo/v/vi200-ic
/usr/lib/terminfo/t/tn1200	/usr/lib/terminfo/v/vi200-rv
/usr/lib/terminfo/t/tn300	/usr/lib/terminfo/v/vi200-rv-ic
/usr/lib/terminfo/t/trs-80	/usr/lib/terminfo/v/vi300
/usr/lib/terminfo/t/trs80	/usr/lib/terminfo/v/vi300-aw
/usr/lib/terminfo/t/ts-1	/usr/lib/terminfo/v/vi300-rv
/usr/lib/terminfo/t/ts1	/usr/lib/terminfo/v/vi300-ss
/usr/lib/terminfo/t/tty	/usr/lib/terminfo/v/vi550
/usr/lib/terminfo/t/tty33	/usr/lib/terminfo/v/viewpoint
/usr/lib/terminfo/t/tty37	/usr/lib/terminfo/v/virtual
/usr/lib/terminfo/t/tty40	/usr/lib/terminfo/v/visual
/usr/lib/terminfo/t/tty43	/usr/lib/terminfo/v/vitty
/usr/lib/terminfo/t/tty4424	/usr/lib/terminfo/v/vk100
/usr/lib/terminfo/t/tty4424-2	/usr/lib/terminfo/v/vs100
/usr/lib/terminfo/t/tty4424-el-2	/usr/lib/terminfo/v/vs100s
/usr/lib/terminfo/t/ttywilliams	/usr/lib/terminfo/v/vt100
/usr/lib/terminfo/t/tvi	/usr/lib/terminfo/v/vt100-am
/usr/lib/terminfo/t/tvi-2p	/usr/lib/terminfo/v/vt100-bot-s
/usr/lib/terminfo/t/tvi2p	/usr/lib/terminfo/v/vt100-nam
/usr/lib/terminfo/v/vt100-nam-w	/usr/lib/terminfo/v/vt220-am
/usr/lib/terminfo/v/vt100-nav	/usr/lib/terminfo/v/vt50
/usr/lib/terminfo/v/vt100-nav-w	/usr/lib/terminfo/v/vt50h
/usr/lib/terminfo/v/vt100-np	/usr/lib/terminfo/v/vt52
/usr/lib/terminfo/v/vt100-s	/usr/lib/terminfo/w/wy100
/usr/lib/terminfo/v/vt100-s-bot	/usr/lib/terminfo/w/wy30
/usr/lib/terminfo/v/vt100-s-top	/usr/lib/terminfo/x/x1700
/usr/lib/terminfo/v/vt100-top-s	/usr/lib/terminfo/x/x1720
/usr/lib/terminfo/v/vt100-w	/usr/lib/terminfo/x/x1750
/usr/lib/terminfo/v/vt100-w-am	/usr/lib/terminfo/x/xitex
/usr/lib/terminfo/v/vt100-w-nam	/usr/lib/terminfo/x/xl83
/usr/lib/terminfo/v/vt100-w-nav	/usr/lib/terminfo/x/xterm

/usr/lib/terminfo/v/vt100am	/usr/lib/terminfo/x/xterms
/usr/lib/terminfo/v/vt100nam	/usr/lib/terminfo/y/ya
/usr/lib/terminfo/v/vt100s	/usr/lib/terminfo/z/z19
/usr/lib/terminfo/v/vt100w	/usr/lib/terminfo/z/z30
/usr/lib/terminfo/v/vt125	/usr/lib/terminfo/z/zen30
/usr/lib/terminfo/v/vt132	/usr/lib/terminfo/z/zenith
/usr/lib/terminfo/v/vt220	

## Fileset LAN

/etc/conf/libatalk65.a	/usr/include/netinet/udp.h
/etc/conf/liblan65.a	/usr/include/netinet/udp_var.h
/system/LAN/customize	/usr/include/sio/iocam.h
/system/LAN/revlist	/usr/include/sio/iodam.h
/usr/include/net/af.h	/usr/include/sio/iodiag.h
/usr/include/net/if.h	/usr/include/sio/iodm.h
/usr/include/net/lan_802.h	/usr/include/sio/llio.h
/usr/include/net/netisr.h	/usr/include/sio/nsdiag0.h
/usr/include/net/probe.h	/usr/include/sys/dctrace.h
/usr/include/net/probe_4.h	/usr/include/sys/domain.h
/usr/include/net/raw8023.h	/usr/include/sys/mbuf.h
/usr/include/net/raw_cb.h	/usr/include/sys/meas_sys.h
/usr/include/net/route.h	/usr/include/sys/narchdef.h
/usr/include/netinet/arp.h	/usr/include/sys/narchvar.h
/usr/include/netinet/icmp_var.h	/usr/include/sys/ndtrace.h
/usr/include/netinet/if_ether.h	/usr/include/sys/nerror.h
/usr/include/netinet/in.h	/usr/include/sys/nmacros.h
/usr/include/netinet/in_pcb.h	/usr/include/sys/ns_canon.h
/usr/include/netinet/in_sysm.h	/usr/include/sys/ns_diag.h
/usr/include/netinet/ip.h	/usr/include/sys/ns_ipc.h
/usr/include/netinet/ip_icmp.h	/usr/include/sys/ns_ipcvar.h
/usr/include/netinet/ip_var.h	/usr/include/sys/ns_names.h
/usr/include/netinet/pmt.h	/usr/include/sys/ns_opt.h
/usr/include/netinet/probe_2.h	/usr/include/sys/ntask.h
/usr/include/netinet/pxp.h	/usr/include/sys/ntypes.h
/usr/include/netinet/pxp_var.h	/usr/include/sys/protosw.h
/usr/include/netinet/tcp.h	/usr/include/sys/socket.h
/usr/include/netinet/tcp_debug.h	/usr/include/sys/socketvar.h
/usr/include/netinet/tcp_fsm.h	/usr/include/sys/syscall.h
/usr/include/netinet/tcp_seq.h	/usr/include/sys/trigdef.h
/usr/include/netinet/tcp_timer.h	/usr/include/sys/un.h
/usr/include/netinet/tcp_var.h	/usr/include/sys/unpcb.h
/usr/include/netinet/tcpip.h	/usr/include/sys/x25_diag.h



## Partition NLS

### Fileset NAMERICA

<code>/system/NAMERICA/revlist</code>	<code>/usr/lib/nls/american/environment</code>
<code>/usr/lib/nls/american/collate8</code>	

### Fileset NCFRENCH

<code>/system/NCFRENCH/revlist</code>	<code>/usr/lib/nls/c-french/environment</code>
<code>/usr/lib/nls/c-french/collate8</code>	

### Fileset NCHINESS

<code>/system/NCHINESS/revlist</code>	<code>/usr/lib/nls/chinese-s/environment</code>
<code>/usr/lib/nls/chinese-s/collate8</code>	

### Fileset NCHINEST

<code>/system/NCHINEST/revlist</code>	<code>/usr/lib/nls/chinese-t/environment</code>
<code>/usr/lib/nls/chinese-t/collate8</code>	

### Fileset NDANISH

<code>/system/NDANISH/revlist</code>	<code>/usr/lib/nls/danish/environment</code>
<code>/usr/lib/nls/danish/collate8</code>	

### Fileset NDUTCH

<code>/system/NDUTCH/revlist</code>	<code>/usr/lib/nls/dutch/environment</code>
<code>/usr/lib/nls/dutch/collate8</code>	

### Fileset NENGLISH

<code>/system/NENGLISH/revlist</code>	<code>/usr/lib/nls/english/environment</code>
<code>/usr/lib/nls/english/collate8</code>	

### Fileset NFINNISH

<code>/system/NFINNISH/revlist</code>	<code>/usr/lib/nls/finnish/environment</code>
<code>/usr/lib/nls/finnish/collate8</code>	

## Fileset NFRENCH

/system/NFRENCH/revlist                   /usr/lib/nls/french/environment  
/usr/lib/nls/french/collate8

## Fileset NGERMAN

/system/NGERMAN/revlist                   /usr/lib/nls/german/environment  
/usr/lib/nls/german/collate8

## Fileset NGREEK

/system/NGREEK/revlist                   /usr/lib/nls/greek/environment  
/usr/lib/nls/greek/collate8

## Fileset NICELAND

/system/NICELAND/revlist                   /usr/lib/nls/icelandic/environment  
/usr/lib/nls/icelandic/collate8

## Fileset NITALIAN

/system/NITALIAN/revlist                   /usr/lib/nls/italian/environment  
/usr/lib/nls/italian/collate8

## Fileset NJAPANES

/system/NJAPANES/customize               /usr/lib/nls/japanese/environment  
/system/NJAPANES/revlist                 /usr/lib/nls/katakana/collate8  
/usr/lib/nls/japanese/collate8            /usr/lib/nls/katakana/environment

## Fileset NKOREAN

/system/NKOREAN/revlist                   /usr/lib/nls/korean/environment  
/usr/lib/nls/korean/collate8

## Fileset NNORWEGI

/system/NNORWEGI/revlist                   /usr/lib/nls/norwegian/environment  
/usr/lib/nls/norwegian/collate8

## Fileset NPORTUGU

/system/NPORTUGU/revlist                   /usr/lib/nls/portuguese/environment  
/usr/lib/nls/portuguese/collate8

## **Fileset NSPANISH**

`/system/NSPANISH/revlist`  
`/usr/lib/nls/spanish/collate8`

`/usr/lib/nls/spanish/environment`

## **Fileset NSWEDISH**

`/system/NSWEDISH/revlist`  
`/usr/lib/nls/swedish/collate8`

`/usr/lib/nls/swedish/environment`

---

## List of PE System Files

The following is a list of the PE partitions and filesets, and the files contained within them (at least the mandatory AXE filesets must be loaded before you can use your PE filesets). This list may vary slightly from what is actually provided on the installation media. Also, some of the files are not actually on the media but are created by the customization scripts. If you want a particular file, load the whole fileset.

### Partition SYS\_CORE

#### Fileset FTN\_CORE

<code>/lib/frt0.o</code>	<code>/usr/lib/libF77.a</code>
<code>/system/FTN_CORE/revlist</code>	<code>/usr/lib/libI77.a</code>

#### Fileset PAS\_CORE

<code>/lib/libpc.a</code>	<code>/usr/lib/libheap2.a</code>
<code>/lib/libpc_p.a</code>	<code>/usr/lib/libheap2_p.a</code>
<code>/system/PAS_CORE/revlist</code>	

#### Fileset PCORE

<code>/bin/crypt</code>	<code>/etc/tunefs</code>
<code>/bin/diff</code>	<code>/etc/whodo</code>
<code>/bin/env</code>	<code>/system/PCORE/CDFcustomize</code>
<code>/bin/mail</code>	<code>/system/PCORE/customize</code>
<code>/bin/mesg</code>	<code>/system/PCORE/revlist</code>
<code>/bin/newgrp</code>	<code>/usr/bin/at</code>
<code>/bin/pty</code>	<code>/usr/bin/batch</code>
<code>/bin/ranlib</code>	<code>/usr/bin/bdiff</code>
<code>/bin/rmail</code>	<code>/usr/bin/bifdf</code>
<code>/bin/sum</code>	<code>/usr/bin/col</code>
<code>/bin/tty</code>	<code>/usr/bin/crontab</code>
<code>/bin/write</code>	<code>/usr/bin/ctags</code>
<code>/etc/clri</code>	<code>/usr/bin/egrep</code>
<code>/etc/cron</code>	<code>/usr/bin/expand</code>
<code>/etc/fsdb</code>	<code>/usr/bin/fgrep</code>
<code>/etc/fuser</code>	<code>/usr/bin/file</code>
<code>/etc/grpck</code>	<code>/usr/bin/getprivgrp</code>
<code>/etc/last</code>	<code>/usr/bin/groups</code>
<code>/etc/lastb</code>	<code>/usr/bin/id</code>

/etc/magic	/usr/bin/logname
/etc/mklost+found	/usr/bin/news
/etc/ncheck	/usr/bin/prealloc
/etc/newconfig/.proto	/usr/bin/rtprio
/etc/newconfig/at.allow	/usr/bin/tput
/etc/newconfig/cron.allow	/usr/bin/unexpand
/etc/newconfig/crontab	/usr/bin/xargs
/etc/newconfig/mklp	/usr/lib/diffh
/etc/newconfig/queuedefs	/usr/lib/help/bd
/etc/newconfig/termcap	/usr/lib/help/term
/etc/ptydaemon	/usr/lib/help/text
/etc/pwck	/usr/lib/makekey
/etc/revck	/usr/lib/nmf
/etc/setprivgrp	

## Fileset PDSKLESS

/bin/cps	/etc/read_cct
/etc/ccck	/system/PDSKLESS/customize
/etc/cfuser	/system/PDSKLESS/revlist
/etc/cluster	/usr/bin/makecdf
/etc/clustersh	/usr/bin/showcdf
/etc/csp	/usr/boot/SYSBKUP
/etc/cwall	/usr/boot/SYSDEBUG
/etc/newconfig/boottab	/usr/boot/SYSHPUX
/etc/rbootd	

## Partition TEXT

### Fileset HPUX\_MAN

/system/HPUX_MAN/customize	/usr/man/cat1.Z/comm.1
/usr/man/cat1.Z/ada.1	/usr/man/cat1.Z/compact.1
/usr/man/cat1.Z/ada.fmgr.1	/usr/man/cat1.Z/compress.1
/usr/man/cat1.Z/ada.format.1	/usr/man/cat1.Z/cp.1
/usr/man/cat1.Z/ada.mklib.1	/usr/man/cat1.Z/cpio.1
/usr/man/cat1.Z/ada.mvlib.1	/usr/man/cat1.Z/cpp.1
/usr/man/cat1.Z/ada.probe.1	/usr/man/cat1.Z/cps.1
/usr/man/cat1.Z/ada.protect.1	/usr/man/cat1.Z/crontab.1
/usr/man/cat1.Z/ada.rmlib.1	/usr/man/cat1.Z/crypt.1
/usr/man/cat1.Z/ada.umgr.1	/usr/man/cat1.Z/csh.1
/usr/man/cat1.Z/ada.unlock.1	/usr/man/cat1.Z/csplit.1
/usr/man/cat1.Z/ada.xref.1	/usr/man/cat1.Z/ct.1
/usr/man/cat1.Z/adb.1	/usr/man/cat1.Z/ctags.1

```

/usr/man/cat1.Z/adjust.1
/usr/man/cat1.Z/admin.1
/usr/man/cat1.Z/ar.1
/usr/man/cat1.Z/as.1
/usr/man/cat1.Z/asa.1
/usr/man/cat1.Z/astrn.1
/usr/man/cat1.Z/at.1
/usr/man/cat1.Z/atime.1
/usr/man/cat1.Z/atrans.1
/usr/man/cat1.Z/awk.1
/usr/man/cat1.Z/banner.1
/usr/man/cat1.Z/basename.1
/usr/man/cat1.Z/basic.1
/usr/man/cat1.Z/batch.1
/usr/man/cat1.Z/bc.1
/usr/man/cat1.Z/bdiff.1
/usr/man/cat1.Z/bfs.1
/usr/man/cat1.Z/bifchgrp.1
/usr/man/cat1.Z/bifchmod.1
/usr/man/cat1.Z/bifchown.1
/usr/man/cat1.Z/bifcp.1
/usr/man/cat1.Z/biffind.1
/usr/man/cat1.Z/bifls.1
/usr/man/cat1.Z/bifmkdir.1
/usr/man/cat1.Z/bifrm.1
/usr/man/cat1.Z/bifrmkdir.1
/usr/man/cat1.Z/bs.1
/usr/man/cat1.Z/cal.1
/usr/man/cat1.Z/calendar.1
/usr/man/cat1.Z/cancel.1
/usr/man/cat1.Z/cu.1
/usr/man/cat1.Z/cut.1
/usr/man/cat1.Z/cxref.1
/usr/man/cat1.Z/date.1
/usr/man/cat1.Z/dc.1
/usr/man/cat1.Z/dd.1
/usr/man/cat1.Z/delta.1
/usr/man/cat1.Z/deroff.1
/usr/man/cat1.Z/diff.1
/usr/man/cat1.Z/diff3.1
/usr/man/cat1.Z/diffh.1
/usr/man/cat1.Z/diffmk.1
/usr/man/cat1.Z/dircmp.1
/usr/man/cat1.Z/dirname.1
/usr/man/cat1.Z/disable.1
/usr/man/cat1.Z/dos2ux.1
/usr/man/cat1.Z/doschmod.1
/usr/man/cat1.Z/doscp.1
/usr/man/cat1.Z/dosdf.1
/usr/man/cat1.Z/dosll.1
/usr/man/cat1.Z/dosls.1
/usr/man/cat1.Z/dosmkdir.1
/usr/man/cat1.Z/dosrm.1
/usr/man/cat1.Z/dosrmdir.1
/usr/man/cat1.Z/du.1
/usr/man/cat1.Z/dumpmsg.1
/usr/man/cat1.Z/echo.1
/usr/man/cat1.Z/ed.1
/usr/man/cat1.Z/edit.1
/usr/man/cat1.Z/egrep.1

```

/usr/man/cat1.Z/cat.1	/usr/man/cat1.Z/enable.1
/usr/man/cat1.Z/cb.1	/usr/man/cat1.Z/env.1
/usr/man/cat1.Z/cc.1	/usr/man/cat1.Z/ex.1
/usr/man/cat1.Z/ccat.1	/usr/man/cat1.Z/expand.1
/usr/man/cat1.Z/cd.1	/usr/man/cat1.Z/expr.1
/usr/man/cat1.Z/cdb.1	/usr/man/cat1.Z/f77.1
/usr/man/cat1.Z/cdc.1	/usr/man/cat1.Z/factor.1
/usr/man/cat1.Z/cflow.1	/usr/man/cat1.Z/false.1
/usr/man/cat1.Z/chacl.1	/usr/man/cat1.Z/fc.1
/usr/man/cat1.Z/chatr.1	/usr/man/cat1.Z/fdb.1
/usr/man/cat1.Z/checknr.1	/usr/man/cat1.Z/fgrep.1
/usr/man/cat1.Z/chfn.1	/usr/man/cat1.Z/file.1
/usr/man/cat1.Z/chgrp.1	/usr/man/cat1.Z/find.1
/usr/man/cat1.Z/chmod.1	/usr/man/cat1.Z/findmsg.1
/usr/man/cat1.Z/chown.1	/usr/man/cat1.Z/findstr.1
/usr/man/cat1.Z/chsh.1	/usr/man/cat1.Z/finger.1
/usr/man/cat1.Z/ci.1	/usr/man/cat1.Z/fixman.1
/usr/man/cat1.Z/clear.1	/usr/man/cat1.Z/flint.1
/usr/man/cat1.Z/cmp.1	/usr/man/cat1.Z/fold.1
/usr/man/cat1.Z/cnodes.1	/usr/man/cat1.Z/forder.1
/usr/man/cat1.Z/co.1	/usr/man/cat1.Z/from.1
/usr/man/cat1.Z/col.1	/usr/man/cat1.Z/fsplit.1
/usr/man/cat1.Z/comb.1	/usr/man/cat1.Z/ftio.1
/usr/man/cat1.Z/gencat.1	/usr/man/cat1.Z/man.1
/usr/man/cat1.Z/get.1	/usr/man/cat1.Z/mediainit.1
/usr/man/cat1.Z/getaccess.1	/usr/man/cat1.Z/merge.1
/usr/man/cat1.Z/getcontext.1	/usr/man/cat1.Z/mesg.1
/usr/man/cat1.Z/getopt.1	/usr/man/cat1.Z/mkdir.1
/usr/man/cat1.Z/getprivgrp.1	/usr/man/cat1.Z/mkmf.1
/usr/man/cat1.Z/gprof.1	/usr/man/cat1.Z/mkstr.1
/usr/man/cat1.Z/grep.1	/usr/man/cat1.Z/mkuupath.1
/usr/man/cat1.Z/grget.1	/usr/man/cat1.Z/mm.1
/usr/man/cat1.Z/groups.1	/usr/man/cat1.Z/more.1
/usr/man/cat1.Z/hashcheck.1	/usr/man/cat1.Z/mt.1
/usr/man/cat1.Z/hashmake.1	/usr/man/cat1.Z/mv.1
/usr/man/cat1.Z/head.1	/usr/man/cat1.Z/neqn.1
/usr/man/cat1.Z/help.1	/usr/man/cat1.Z/newform.1
/usr/man/cat1.Z/hostname.1	/usr/man/cat1.Z/newgrp.1
/usr/man/cat1.Z/hp.1	/usr/man/cat1.Z/news.1
/usr/man/cat1.Z/hp9000s200.1	/usr/man/cat1.Z/nice.1
/usr/man/cat1.Z/hp9000s300.1	/usr/man/cat1.Z/nl.1
/usr/man/cat1.Z/hp9000s500.1	/usr/man/cat1.Z/nlio.1
/usr/man/cat1.Z/hp9000s800.1	/usr/man/cat1.Z/nliov.1

/usr/man/cat1.Z/hyphen.1	/usr/man/cat1.Z/nliostart.1
/usr/man/cat1.Z/id.1	/usr/man/cat1.Z/nljust.1
/usr/man/cat1.Z/ident.1	/usr/man/cat1.Z/nm.1
/usr/man/cat1.Z/insertmsg.1	/usr/man/cat1.Z/nohup.1
/usr/man/cat1.Z/intro.1	/usr/man/cat1.Z/nroff.1
/usr/man/cat1.Z/inv.1	/usr/man/cat1.Z/od.1
/usr/man/cat1.Z/ipcrm.1	/usr/man/cat1.Z/pack.1
/usr/man/cat1.Z/ipcs.1	/usr/man/cat1.Z/page.1
/usr/man/cat1.Z/isql.1	/usr/man/cat1.Z/pam.1
/usr/man/cat1.Z/join.1	/usr/man/cat1.Z/passwd.1
/usr/man/cat1.Z/kermit.1	/usr/man/cat1.Z/paste.1
/usr/man/cat1.Z/kill.1	/usr/man/cat1.Z/pathalias.1
/usr/man/cat1.Z/ksh.1	/usr/man/cat1.Z/pc.1
/usr/man/cat1.Z/l.1	/usr/man/cat1.Z/pcat.1
/usr/man/cat1.Z/last.1	/usr/man/cat1.Z/pdb.1
/usr/man/cat1.Z/lastb.1	/usr/man/cat1.Z/pdp11.1
/usr/man/cat1.Z/lastcomm.1	/usr/man/cat1.Z/pg.1
/usr/man/cat1.Z/ld.1	/usr/man/cat1.Z/pr.1
/usr/man/cat1.Z/leave.1	/usr/man/cat1.Z/prealloc.1
/usr/man/cat1.Z/lex.1	/usr/man/cat1.Z/primes.1
/usr/man/cat1.Z/lifcp.1	/usr/man/cat1.Z/printenv.1
/usr/man/cat1.Z/lifinit.1	/usr/man/cat1.Z/prmail.1
/usr/man/cat1.Z/lifls.1	/usr/man/cat1.Z/prof.1
/usr/man/cat1.Z/lifrename.1	/usr/man/cat1.Z/prs.1
/usr/man/cat1.Z/lifrm.1	/usr/man/cat1.Z/ps.1
/usr/man/cat1.Z/line.1	/usr/man/cat1.Z/psqlc.1
/usr/man/cat1.Z/lint.1	/usr/man/cat1.Z/psqlcbl.1
/usr/man/cat1.Z/ll.1	/usr/man/cat1.Z/psqlfor.1
/usr/man/cat1.Z/ln.1	/usr/man/cat1.Z/psqlpas.1
/usr/man/cat1.Z/lock.1	/usr/man/cat1.Z/ptx.1
/usr/man/cat1.Z/login.1	/usr/man/cat1.Z/pty.1
/usr/man/cat1.Z/logname.1	/usr/man/cat1.Z/pwd.1
/usr/man/cat1.Z/lorder.1	/usr/man/cat1.Z/pwget.1
/usr/man/cat1.Z/lp.1	/usr/man/cat1.Z/ratfor.1
/usr/man/cat1.Z/lpstat.1	/usr/man/cat1.Z/rcs.1
/usr/man/cat1.Z/lrs.1	/usr/man/cat1.Z/rcsdiff.1
/usr/man/cat1.Z/lrsacl.1	/usr/man/cat1.Z/rcsmerge.1
/usr/man/cat1.Z/lrsf.1	/usr/man/cat1.Z/red.1
/usr/man/cat1.Z/lrsr.1	/usr/man/cat1.Z/rev.1
/usr/man/cat1.Z/lrsx.1	/usr/man/cat1.Z/rksh.1
/usr/man/cat1.Z/m4.1	/usr/man/cat1.Z/rlog.1
/usr/man/cat1.Z/machid.1	/usr/man/cat1.Z/rm.1



/usr/man/cat1.Z/mail.1	/usr/man/cat1.Z/rmail.1
/usr/man/cat1.Z/mailx.1	/usr/man/cat1.Z/rmdel.1
/usr/man/cat1.Z/make.1	/usr/man/cat1.Z/rmdir.1
/usr/man/cat1.Z/makekey.1	/usr/man/cat1.Z/rmnl.1
/usr/man/cat1.Z/rsh.1	/usr/man/cat1.Z/uptime.1
/usr/man/cat1.Z/rtprio.1	/usr/man/cat1.Z/users.1
/usr/man/cat1.Z/sact.1	/usr/man/cat1.Z/uucp.1
/usr/man/cat1.Z/sccsdiff.1	/usr/man/cat1.Z/uulog.1
/usr/man/cat1.Z/script.1	/usr/man/cat1.Z/uuname.1
/usr/man/cat1.Z/sdfchgrp.1	/usr/man/cat1.Z/uupath.1
/usr/man/cat1.Z/sdfchmod.1	/usr/man/cat1.Z/uupick.1
/usr/man/cat1.Z/sdfchown.1	/usr/man/cat1.Z/uustat.1
/usr/man/cat1.Z/sdfcp.1	/usr/man/cat1.Z/uuto.1
/usr/man/cat1.Z/sdffind.1	/usr/man/cat1.Z/uux.1
/usr/man/cat1.Z/sdfll.1	/usr/man/cat1.Z/ux2dos.1
/usr/man/cat1.Z/sdfln.1	/usr/man/cat1.Z/val.1
/usr/man/cat1.Z/sdfls.1	/usr/man/cat1.Z/vax.1
/usr/man/cat1.Z/sdfmkdir.1	/usr/man/cat1.Z/vc.1
/usr/man/cat1.Z/sdfmv.1	/usr/man/cat1.Z/vi.1
/usr/man/cat1.Z/sdfrm.1	/usr/man/cat1.Z/vis.1
/usr/man/cat1.Z/sdfrmdir.1	/usr/man/cat1.Z/vt.1
/usr/man/cat1.Z/sdiff.1	/usr/man/cat1.Z/wait.1
/usr/man/cat1.Z/sed.1	/usr/man/cat1.Z/wc.1
/usr/man/cat1.Z/sh.1	/usr/man/cat1.Z/wdedit.1
/usr/man/cat1.Z/shar.1	/usr/man/cat1.Z/wdutil.1
/usr/man/cat1.Z/shl.1	/usr/man/cat1.Z/what.1
/usr/man/cat1.Z/showcdf.1	/usr/man/cat1.Z/whereis.1
/usr/man/cat1.Z/size.1	/usr/man/cat1.Z/which.1
/usr/man/cat1.Z/sleep.1	/usr/man/cat1.Z/who.1
/usr/man/cat1.Z/slp.1	/usr/man/cat1.Z/whoami.1
/usr/man/cat1.Z/soelim.1	/usr/man/cat1.Z/write.1
/usr/man/cat1.Z/sort.1	/usr/man/cat1.Z/xargs.1
/usr/man/cat1.Z/spell.1	/usr/man/cat1.Z/xd.1
/usr/man/cat1.Z/spellin.1	/usr/man/cat1.Z/xdb.1
/usr/man/cat1.Z/split.1	/usr/man/cat1.Z/yacc.1
/usr/man/cat1.Z/sqlutil.1	/usr/man/cat1.Z/yes.1
/usr/man/cat1.Z/ssp.1	/usr/man/cat1.Z/zcat.1
/usr/man/cat1.Z/strings.1	/usr/man/cat1m.Z/accept.1m
/usr/man/cat1.Z/strip.1	/usr/man/cat1m.Z/acct.1m
/usr/man/cat1.Z/stty.1	/usr/man/cat1m.Z/acctcms.1m
/usr/man/cat1.Z/su.1	/usr/man/cat1m.Z/acctcom.1m
/usr/man/cat1.Z/sum.1	/usr/man/cat1m.Z/acctcon.1m

```

/usr/man/cat1.Z/tabs.1
/usr/man/cat1.Z/tail.1
/usr/man/cat1.Z/tar.1
/usr/man/cat1.Z/tbl.1
/usr/man/cat1.Z/tcio.1
/usr/man/cat1.Z/tee.1
/usr/man/cat1.Z/test.1
/usr/man/cat1.Z/time.1
/usr/man/cat1.Z/touch.1
/usr/man/cat1.Z/tput.1
/usr/man/cat1.Z/tr.1
/usr/man/cat1.Z/true.1
/usr/man/cat1.Z/tset.1
/usr/man/cat1.Z/tsort.1
/usr/man/cat1.Z/tty.1
/usr/man/cat1.Z/u3b.1
/usr/man/cat1.Z/u3b5.1
/usr/man/cat1.Z/ul.1
/usr/man/cat1.Z/umask.1
/usr/man/cat1.Z/umodem.1
/usr/man/cat1.Z/uname.1
/usr/man/cat1.Z/uncompact.1
/usr/man/cat1.Z/uncompress.1
/usr/man/cat1.Z/unexpand.1
/usr/man/cat1.Z/unget.1
/usr/man/cat1.Z/uniq.1
/usr/man/cat1.Z/units.1
/usr/man/cat1.Z/unpack.1
/usr/man/cat1m.Z/chargefee.1m
/usr/man/cat1m.Z/chroot.1m
/usr/man/cat1m.Z/ckpacct.1m
/usr/man/cat1m.Z/clri.1m
/usr/man/cat1m.Z/clrsvc.1m
/usr/man/cat1m.Z/cluster.1m
/usr/man/cat1m.Z/config.1m
/usr/man/cat1m.Z/convertfs.1m
/usr/man/cat1m.Z/cpset.1m
/usr/man/cat1m.Z/cron.1m
/usr/man/cat1m.Z/csp.1m
/usr/man/cat1m.Z/cwall.1m
/usr/man/cat1m.Z/devnm.1m
/usr/man/cat1m.Z/df.1m
/usr/man/cat1m.Z/acctcon1.1m
/usr/man/cat1m.Z/acctcon2.1m
/usr/man/cat1m.Z/acctdisk.1m
/usr/man/cat1m.Z/acctdusg.1m
/usr/man/cat1m.Z/acctmerg.1m
/usr/man/cat1m.Z/accton.1m
/usr/man/cat1m.Z/acctprc.1m
/usr/man/cat1m.Z/acctprc1.1m
/usr/man/cat1m.Z/acctprc2.1m
/usr/man/cat1m.Z/acctsh.1m
/usr/man/cat1m.Z/acctwtmp.1m
/usr/man/cat1m.Z/audevent.1m
/usr/man/cat1m.Z/audisp.1m
/usr/man/cat1m.Z/audomon.1m
/usr/man/cat1m.Z/audsys.1m
/usr/man/cat1m.Z/audusr.1m
/usr/man/cat1m.Z/backup.1m
/usr/man/cat1m.Z/bcheckrc.1m
/usr/man/cat1m.Z/bdf.1m
/usr/man/cat1m.Z/bifdf.1m
/usr/man/cat1m.Z/biffsc.1m
/usr/man/cat1m.Z/biffscdb.1m
/usr/man/cat1m.Z/bifmkfs.1m
/usr/man/cat1m.Z/brc.1m
/usr/man/cat1m.Z/captoinfo.1m
/usr/man/cat1m.Z/catman.1m
/usr/man/cat1m.Z/ccck.1m
/usr/man/cat1m.Z/cfuser.1m
/usr/man/cat1m.Z/reject.1m
/usr/man/cat1m.Z/restore.1m
/usr/man/cat1m.Z/revck.1m
/usr/man/cat1m.Z/rlp.1m
/usr/man/cat1m.Z/rlpdaemon.1m
/usr/man/cat1m.Z/rlpstat.1m
/usr/man/cat1m.Z/runacct.1m
/usr/man/cat1m.Z/sdfdf.1m
/usr/man/cat1m.Z/sdffsc.1m
/usr/man/cat1m.Z/sdffscdb.1m
/usr/man/cat1m.Z/setmnt.1m
/usr/man/cat1m.Z/setprivgrp.1m
/usr/man/cat1m.Z/shutacct.1m
/usr/man/cat1m.Z/shutdown.1m

```

```

/usr/man/cat1m.Z/diskinfo.1m
/usr/man/cat1m.Z/diskusg.1m
/usr/man/cat1m.Z/dmesg.1m
/usr/man/cat1m.Z/dodisk.1m
/usr/man/cat1m.Z/dump.1m
/usr/man/cat1m.Z/fbackup.1m
/usr/man/cat1m.Z/frecover.1m
/usr/man/cat1m.Z/fsck.1m
/usr/man/cat1m.Z/fsclean.1m
/usr/man/cat1m.Z/fsdb.1m
/usr/man/cat1m.Z/fstomnt.1m
/usr/man/cat1m.Z/fuser.1m
/usr/man/cat1m.Z/fwtmp.1m
/usr/man/cat1m.Z/getty.1m
/usr/man/cat1m.Z/getx25.1m
/usr/man/cat1m.Z/grpck.1m
/usr/man/cat1m.Z/init.1m
/usr/man/cat1m.Z/install.1m
/usr/man/cat1m.Z/intro.1m
/usr/man/cat1m.Z/killall.1m
/usr/man/cat1m.Z/lastlogin.1m
/usr/man/cat1m.Z/link.1m
/usr/man/cat1m.Z/lpadmin.1m
/usr/man/cat1m.Z/lpmove.1m
/usr/man/cat1m.Z/lpsched.1m
/usr/man/cat1m.Z/lpshut.1m
/usr/man/cat1m.Z/lsdev.1m
/usr/man/cat1m.Z/makecdf.1m
/usr/man/cat1m.Z/mkdev.1m
/usr/man/cat1m.Z/mkfs.1m
/usr/man/cat1m.Z/mklost+foun.1m
/usr/man/cat1m.Z/mklp.1m
/usr/man/cat1m.Z/mknod.1m
/usr/man/cat1m.Z/mkrs.1m
/usr/man/cat1m.Z/monacct.1m
/usr/man/cat1m.Z/mount.1m
/usr/man/cat1m.Z/mvdir.1m
/usr/man/cat1m.Z/ncheck.1m
/usr/man/cat1m.Z/newfs.1m
/usr/man/cat1m.Z/nlioinit.1m
/usr/man/cat1m.Z/nulladm.1m
/usr/man/cat1m.Z/opx25.1m
/usr/man/cat1m.Z/startup.1m
/usr/man/cat1m.Z/swapon.1m
/usr/man/cat1m.Z/sync.1m
/usr/man/cat1m.Z/syncer.1m
/usr/man/cat1m.Z/sysrm.1m
/usr/man/cat1m.Z/telinit.1m
/usr/man/cat1m.Z/tic.1m
/usr/man/cat1m.Z/tunefs.1m
/usr/man/cat1m.Z/turnacct.1m
/usr/man/cat1m.Z/umount.1m
/usr/man/cat1m.Z/unlink.1m
/usr/man/cat1m.Z/untic.1m
/usr/man/cat1m.Z/update.1m
/usr/man/cat1m.Z/uuccheck.1m
/usr/man/cat1m.Z/uucico.1m
/usr/man/cat1m.Z/uuclean.1m
/usr/man/cat1m.Z/uucleanup.1m
/usr/man/cat1m.Z/uugetty.1m
/usr/man/cat1m.Z/uuls.1m
/usr/man/cat1m.Z/uusched.1m
/usr/man/cat1m.Z/uusnap.1m
/usr/man/cat1m.Z/uusnaps.1m
/usr/man/cat1m.Z/uusub.1m
/usr/man/cat1m.Z/uuxqt.1m
/usr/man/cat1m.Z/vipw.1m
/usr/man/cat1m.Z/vtdaemon.1m
/usr/man/cat1m.Z/wall.1m
/usr/man/cat1m.Z/whodo.1m
/usr/man/cat1m.Z/wtmpfix.1m
/usr/man/cat2.Z/_exit.2
/usr/man/cat2.Z/access.2
/usr/man/cat2.Z/acct.2
/usr/man/cat2.Z/alarm.2
/usr/man/cat2.Z/audctl.2
/usr/man/cat2.Z/audswtch.2
/usr/man/cat2.Z/audwrite.2
/usr/man/cat2.Z/brk.2
/usr/man/cat2.Z/bsdproc.2
/usr/man/cat2.Z/chdir.2
/usr/man/cat2.Z/chmod.2
/usr/man/cat2.Z/chown.2
/usr/man/cat2.Z/chroot.2

```

/usr/man/cat1m.Z/powerfail.1m	/usr/man/cat2.Z/close.2
/usr/man/cat1m.Z/prctmp.1m	/usr/man/cat2.Z/cnodeid.2
/usr/man/cat1m.Z/prdaily.1m	/usr/man/cat2.Z/cnodes.2
/usr/man/cat1m.Z/prtacct.1m	/usr/man/cat2.Z/creat.2
/usr/man/cat1m.Z/pwck.1m	/usr/man/cat2.Z/dup.2
/usr/man/cat1m.Z/rbootd.1m	/usr/man/cat2.Z/dup2.2
/usr/man/cat1m.Z/rc.1m	/usr/man/cat2.Z/errno.2
/usr/man/cat1m.Z/rcancel.1m	/usr/man/cat2.Z/exec.2
/usr/man/cat1m.Z/reboot.1m	/usr/man/cat2.Z/execl.2
/usr/man/cat1m.Z/reconfig.1m	/usr/man/cat2.Z/execle.2
/usr/man/cat2.Z/execlp.2	/usr/man/cat2.Z/ptrace.2
/usr/man/cat2.Z/execv.2	/usr/man/cat2.Z/read.2
/usr/man/cat2.Z/execve.2	/usr/man/cat2.Z/readlink.2
/usr/man/cat2.Z/execvp.2	/usr/man/cat2.Z/readv.2
/usr/man/cat2.Z/exit.2	/usr/man/cat2.Z/reboot.2
/usr/man/cat2.Z/fchmod.2	/usr/man/cat2.Z/rename.2
/usr/man/cat2.Z/fchown.2	/usr/man/cat2.Z/rmdir.2
/usr/man/cat2.Z/fcntl.2	/usr/man/cat2.Z/rtprio.2
/usr/man/cat2.Z/fgetacl.2	/usr/man/cat2.Z/sbrk.2
/usr/man/cat2.Z/fgetacl.2	/usr/man/cat2.Z/select.2
/usr/man/cat2.Z/fork.2	/usr/man/cat2.Z/semctl.2
/usr/man/cat2.Z/fpathconf.2	/usr/man/cat2.Z/semget.2
/usr/man/cat2.Z/fsetacl.2	/usr/man/cat2.Z/semop.2
/usr/man/cat2.Z/fsetacl.2	/usr/man/cat2.Z/setacl.2
/usr/man/cat2.Z/fstat.2	/usr/man/cat2.Z/setaudit.2
/usr/man/cat2.Z/fstatfs.2	/usr/man/cat2.Z/setauditproc.2
/usr/man/cat2.Z/ftsync.2	/usr/man/cat2.Z/setevent.2
/usr/man/cat2.Z/ftime.2	/usr/man/cat2.Z/setgid.2
/usr/man/cat2.Z/ftruncate.2	/usr/man/cat2.Z/setgroups.2
/usr/man/cat2.Z/getaccess.2	/usr/man/cat2.Z/sethostname.2
/usr/man/cat2.Z/getacl.2	/usr/man/cat2.Z/setitimer.2
/usr/man/cat2.Z/getaudit.2	/usr/man/cat2.Z/setpgrp.2
/usr/man/cat2.Z/getauditproc.2	/usr/man/cat2.Z/setpgrp2.2
/usr/man/cat2.Z/getcontext.2	/usr/man/cat2.Z/setprigrp.2
/usr/man/cat2.Z/getdirentri.2	/usr/man/cat2.Z/setresgid.2
/usr/man/cat2.Z/getegid.2	/usr/man/cat2.Z/setresuid.2
/usr/man/cat2.Z/geteuid.2	/usr/man/cat2.Z/settimeofday.2
/usr/man/cat2.Z/getevent.2	/usr/man/cat2.Z/setuid.2
/usr/man/cat2.Z/getgid.2	/usr/man/cat2.Z/shmat.2
/usr/man/cat2.Z/getgroups.2	/usr/man/cat2.Z/shmctl.2
/usr/man/cat2.Z/gethostname.2	/usr/man/cat2.Z/shmdt.2
/usr/man/cat2.Z/getitimer.2	/usr/man/cat2.Z/shmget.2

```

/usr/man/cat2.Z/getpgrp.2
/usr/man/cat2.Z/getpgrp2.2
/usr/man/cat2.Z/getpid.2
/usr/man/cat2.Z/getppid.2
/usr/man/cat2.Z/getprivgrp.2
/usr/man/cat2.Z/gettimeofda.2
/usr/man/cat2.Z/getuid.2
/usr/man/cat2.Z/gtty.2
/usr/man/cat2.Z/intro.2
/usr/man/cat2.Z/ioctl.2
/usr/man/cat2.Z/kill.2
/usr/man/cat2.Z/killpg.2
/usr/man/cat2.Z/link.2
/usr/man/cat2.Z/lockf.2
/usr/man/cat2.Z/lseek.2
/usr/man/cat2.Z/lstat.2
/usr/man/cat2.Z/lsync.2
/usr/man/cat2.Z/mkdir.2
/usr/man/cat2.Z/mknod.2
/usr/man/cat2.Z/mkrnod.2
/usr/man/cat2.Z/mount.2
/usr/man/cat2.Z/msgctl.2
/usr/man/cat2.Z/msgget.2
/usr/man/cat2.Z/msgop.2
/usr/man/cat2.Z/msgrcv.2
/usr/man/cat2.Z/msgsnd.2
/usr/man/cat2.Z/nice.2
/usr/man/cat2.Z/open.2
/usr/man/cat2.Z/pathconf.2
/usr/man/cat2.Z/pause.2
/usr/man/cat2.Z/pipe.2
/usr/man/cat2.Z/plock.2
/usr/man/cat2.Z/prealloc.2
/usr/man/cat2.Z/profil.2
/usr/man/cat3.Z/CHARADV.3c
/usr/man/cat3.Z/CHARAT.3c
/usr/man/cat3.Z/ERROR.3x
/usr/man/cat3.Z/FIRSTof2.3c
/usr/man/cat3.Z/GETC.3x
/usr/man/cat3.Z/INIT.3x
/usr/man/cat3.Z/PCHAR.3c
/usr/man/cat3.Z/PCHARADV.3c
/usr/man/cat2.Z/shmop.2
/usr/man/cat2.Z/sigblock.2
/usr/man/cat2.Z/signal.2
/usr/man/cat2.Z/sigpause.2
/usr/man/cat2.Z/sigsetmask.2
/usr/man/cat2.Z/sigsetspace.2
/usr/man/cat2.Z/sigvec.2
/usr/man/cat2.Z/sigvector.2
/usr/man/cat2.Z/stat.2
/usr/man/cat2.Z/statfs.2
/usr/man/cat2.Z/stime.2
/usr/man/cat2.Z/stty.2
/usr/man/cat2.Z/swapon.2
/usr/man/cat2.Z/symlink.2
/usr/man/cat2.Z/sync.2
/usr/man/cat2.Z/sysconf.2
/usr/man/cat2.Z/time.2
/usr/man/cat2.Z/times.2
/usr/man/cat2.Z/truncate.2
/usr/man/cat2.Z/ulimit.2
/usr/man/cat2.Z/umask.2
/usr/man/cat2.Z/umount.2
/usr/man/cat2.Z/uname.2
/usr/man/cat2.Z/unlink.2
/usr/man/cat2.Z/ustat.2
/usr/man/cat2.Z/utime.2
/usr/man/cat2.Z/vfork.2
/usr/man/cat2.Z/vfsmount.2
/usr/man/cat2.Z/wait.2
/usr/man/cat2.Z/wait3.2
/usr/man/cat2.Z/write.2
/usr/man/cat2.Z/writev.2
/usr/man/cat3.Z/ADVANCE.3c
/usr/man/cat3.Z/BYTE_STATUS.3c
/usr/man/cat3.Z/dbm.3x
/usr/man/cat3.Z/dbm_clearer.3x
/usr/man/cat3.Z/dbm_close.3x
/usr/man/cat3.Z/dbm_delete.3x
/usr/man/cat3.Z/dbm_error.3x
/usr/man/cat3.Z/dbm_fetch.3x
/usr/man/cat3.Z/dbm_firstke.3x
/usr/man/cat3.Z/dbm_nextkey.3x

```

```

/usr/man/cat3.Z/PEEKC.3x      /usr/man/cat3.Z/dbm_open.3x
/usr/man/cat3.Z/RETURN.3x     /usr/man/cat3.Z/dbm_store.3x
/usr/man/cat3.Z/SECOf2.3c     /usr/man/cat3.Z/dbmclose.3x
/usr/man/cat3.Z/UNGETC.3x    /usr/man/cat3.Z/dbminit.3x
/usr/man/cat3.Z/WCHAR.3c     /usr/man/cat3.Z/delete.3x
/usr/man/cat3.Z/WCHARADV.3c  /usr/man/cat3.Z/dial.3c
/usr/man/cat3.Z/_tolower.3c   /usr/man/cat3.Z/directory.3c
/usr/man/cat3.Z/_toupper.3c   /usr/man/cat3.Z/drand48.3c
/usr/man/cat3.Z/a64l.3c       /usr/man/cat3.Z/ecvt.3c
/usr/man/cat3.Z/abort.3c      /usr/man/cat3.Z/edata.3c
/usr/man/cat3.Z/abs.3c        /usr/man/cat3.Z/encrypt.3c
/usr/man/cat3.Z/acltostr.3c   /usr/man/cat3.Z/end.3c
/usr/man/cat3.Z/acos.3m       /usr/man/cat3.Z/endccent.3c
/usr/man/cat3.Z/addmntent.3x  /usr/man/cat3.Z/endsent.3x
/usr/man/cat3.Z/advance.3x    /usr/man/cat3.Z/endlrent.3c
/usr/man/cat3.Z/almanac.3x    /usr/man/cat3.Z/endlmntent.3x
/usr/man/cat3.Z/asctime.3c    /usr/man/cat3.Z/endpwent.3c
/usr/man/cat3.Z/asin.3m       /usr/man/cat3.Z/endutent.3c
/usr/man/cat3.Z/assert.3x     /usr/man/cat3.Z/erand48.3c
/usr/man/cat3.Z/atan.3m       /usr/man/cat3.Z/erf.3m
/usr/man/cat3.Z/atan2.3m      /usr/man/cat3.Z/erfc.3m
/usr/man/cat3.Z/atof.3c       /usr/man/cat3.Z/errno.3c
/usr/man/cat3.Z/atoi.3c      /usr/man/cat3.Z/etext.3c
/usr/man/cat3.Z/atol.3c       /usr/man/cat3.Z/exp.3m
/usr/man/cat3.Z/bessel.3m     /usr/man/cat3.Z/fabs.3m
/usr/man/cat3.Z/bsearch.3c    /usr/man/cat3.Z/fclose.3s
/usr/man/cat3.Z/byte_status.3c /usr/man/cat3.Z/fcpacl.3c
/usr/man/cat3.Z/calendar.3x   /usr/man/cat3.Z/fcpacl.3c
/usr/man/cat3.Z/calloc.3c     /usr/man/cat3.Z/fcvt.3c
/usr/man/cat3.Z/calloc.3x     /usr/man/cat3.Z/fdopen.3s
/usr/man/cat3.Z/catclose.3c    /usr/man/cat3.Z/feof.3s
/usr/man/cat3.Z/catgetmsg.3c   /usr/man/cat3.Z/ferror.3s
/usr/man/cat3.Z/catgets.3c    /usr/man/cat3.Z/fetch.3x
/usr/man/cat3.Z/catopen.3c    /usr/man/cat3.Z/fflush.3s
/usr/man/cat3.Z/catread.3c    /usr/man/cat3.Z/fgetc.3s
/usr/man/cat3.Z/ceil.3m       /usr/man/cat3.Z/fgetcct.3c
/usr/man/cat3.Z/chownacl.3c   /usr/man/cat3.Z/fgetgrent.3c
/usr/man/cat3.Z/clearerr.3s   /usr/man/cat3.Z/fgetpwent.3c
/usr/man/cat3.Z/clock.3c      /usr/man/cat3.Z/fgets.3s
/usr/man/cat3.Z/clock.3x      /usr/man/cat3.Z/fileno.3s
/usr/man/cat3.Z/closedir.3c   /usr/man/cat3.Z/firstkey.3x
/usr/man/cat3.Z/compile.3x    /usr/man/cat3.Z/firstof2.3c

```

/usr/man/cat3.Z/conv.3c	/usr/man/cat3.Z/floor.3m
/usr/man/cat3.Z/cos.3m	/usr/man/cat3.Z/fmod.3m
/usr/man/cat3.Z/cosh.3m	/usr/man/cat3.Z/fopen.3s
/usr/man/cat3.Z/cpacl.3c	/usr/man/cat3.Z/fprintf.3s
/usr/man/cat3.Z/crt0.3	/usr/man/cat3.Z/fprintmsg.3c
/usr/man/cat3.Z/crt0.o.3	/usr/man/cat3.Z/fputc.3s
/usr/man/cat3.Z/crypt.3c	/usr/man/cat3.Z/fputs.3s
/usr/man/cat3.Z/ctermid.3s	/usr/man/cat3.Z/fread.3s
/usr/man/cat3.Z/ctime.3c	/usr/man/cat3.Z/free.3c
/usr/man/cat3.Z/ctype.3c	/usr/man/cat3.Z/free.3x
/usr/man/cat3.Z/currlangid.3c	/usr/man/cat3.Z/freopen.3s
/usr/man/cat3.Z/curses.3x	/usr/man/cat3.Z/frexp.3c
/usr/man/cat3.Z/cuserid.3s	/usr/man/cat3.Z/frt0.o.3
/usr/man/cat3.Z/cvtnum.3c	/usr/man/cat3.Z/fscanf.3s
/usr/man/cat3.Z/datalock.3c	/usr/man/cat3.Z/fseek.3s
/usr/man/cat3.Z/daylight.3c	/usr/man/cat3.Z/fsetaclentr.3c
/usr/man/cat3.Z/ftell.3s	/usr/man/cat3.Z/io_get_term.3i
/usr/man/cat3.Z/ftw.3c	/usr/man/cat3.Z/io_interrup.3i
/usr/man/cat3.Z/ftwh.3c	/usr/man/cat3.Z/io_lock.3i
/usr/man/cat3.Z/fwrite.3s	/usr/man/cat3.Z/io_on_inter.3i
/usr/man/cat3.Z/gamma.3m	/usr/man/cat3.Z/io_reset.3i
/usr/man/cat3.Z/gcvt.3c	/usr/man/cat3.Z/io_speed_ct.3i
/usr/man/cat3.Z/getc.3s	/usr/man/cat3.Z/io_timeout_.3i
/usr/man/cat3.Z/getccid.3c	/usr/man/cat3.Z/io_unlock.3i
/usr/man/cat3.Z/getccent.3c	/usr/man/cat3.Z/io_width_ct.3i
/usr/man/cat3.Z/getccnam.3c	/usr/man/cat3.Z/is_68010_pr.3c
/usr/man/cat3.Z/getcdf.3c	/usr/man/cat3.Z/is_68881_pr.3c
/usr/man/cat3.Z/getchar.3s	/usr/man/cat3.Z/is_98248A_p.3c
/usr/man/cat3.Z/getcwd.3c	/usr/man/cat3.Z/is_98635A_p.3c
/usr/man/cat3.Z/getenv.3c	/usr/man/cat3.Z/is_hw_prese.3c
/usr/man/cat3.Z/getfsent.3x	/usr/man/cat3.Z/isalnum.3c
/usr/man/cat3.Z/getfsfile.3x	/usr/man/cat3.Z/isalpha.3c
/usr/man/cat3.Z/getfsspec.3x	/usr/man/cat3.Z/isascii.3c
/usr/man/cat3.Z/getfstype.3x	/usr/man/cat3.Z/isatty.3c
/usr/man/cat3.Z/getgrent.3c	/usr/man/cat3.Z/iscntrl.3c
/usr/man/cat3.Z/getgrgid.3c	/usr/man/cat3.Z/isdigit.3c
/usr/man/cat3.Z/getgrnam.3c	/usr/man/cat3.Z/isgraph.3c
/usr/man/cat3.Z/gethcwd.3c	/usr/man/cat3.Z/islower.3c
/usr/man/cat3.Z/getlogin.3c	/usr/man/cat3.Z/isprint.3c
/usr/man/cat3.Z/getmntent.3x	/usr/man/cat3.Z/ispunct.3c
/usr/man/cat3.Z/getmsg.3c	/usr/man/cat3.Z/isspace.3c
/usr/man/cat3.Z/getopt.3c	/usr/man/cat3.Z/isupper.3c

/usr/man/cat3.Z/getpass.3c	/usr/man/cat3.Z/isxdigit.3c
/usr/man/cat3.Z/getpw.3c	/usr/man/cat3.Z/j0.3m
/usr/man/cat3.Z/getpwent.3c	/usr/man/cat3.Z/j1.3m
/usr/man/cat3.Z/getpwnam.3c	/usr/man/cat3.Z/jn.3m
/usr/man/cat3.Z/getpwuid.3c	/usr/man/cat3.Z/jrand48.3c
/usr/man/cat3.Z/gets.3s	/usr/man/cat3.Z/l3tol.3c
/usr/man/cat3.Z/getspwent.3c	/usr/man/cat3.Z/l64a.3c
/usr/man/cat3.Z/getut.3c	/usr/man/cat3.Z/langinfo.3c
/usr/man/cat3.Z/getutent.3c	/usr/man/cat3.Z/langinit.3c
/usr/man/cat3.Z/getutid.3c	/usr/man/cat3.Z/langtoid.3c
/usr/man/cat3.Z/getutline.3c	/usr/man/cat3.Z/lcong48.3c
/usr/man/cat3.Z/getw.3s	/usr/man/cat3.Z/ldexp.3c
/usr/man/cat3.Z/gmtime.3c	/usr/man/cat3.Z/lfind.3c
/usr/man/cat3.Z/gpio_get_st.3i	/usr/man/cat3.Z/localtime.3c
/usr/man/cat3.Z/gpio_set_ct.3i	/usr/man/cat3.Z/log.3m
/usr/man/cat3.Z/gsignal.3c	/usr/man/cat3.Z/log10.3m
/usr/man/cat3.Z/hasmntopt.3x	/usr/man/cat3.Z/logname.3c
/usr/man/cat3.Z/hcreate.3c	/usr/man/cat3.Z/longjmp.3c
/usr/man/cat3.Z/hdestroy.3c	/usr/man/cat3.Z/lrand48.3c
/usr/man/cat3.Z/hpib_abort.3i	/usr/man/cat3.Z/lsearch.3c
/usr/man/cat3.Z/hpib_bus_st.3i	/usr/man/cat3.Z/ltol3.3c
/usr/man/cat3.Z/hpib_card_p.3i	/usr/man/cat3.Z/mallinfo.3x
/usr/man/cat3.Z/hpib_eoi_ct.3i	/usr/man/cat3.Z/malloc.3c
/usr/man/cat3.Z/hpib_io.3i	/usr/man/cat3.Z/malloc.3x
/usr/man/cat3.Z/hpib_pass_c.3i	/usr/man/cat3.Z/mallopt.3x
/usr/man/cat3.Z/hpib_ppoll.3i	/usr/man/cat3.Z/matherr.3m
/usr/man/cat3.Z/hpib_ppoll_.3i	/usr/man/cat3.Z/mcrt0.o.3
/usr/man/cat3.Z/hpib_ren_ct.3i	/usr/man/cat3.Z/memccpy.3c
/usr/man/cat3.Z/hpib_rqst_s.3i	/usr/man/cat3.Z/memchr.3c
/usr/man/cat3.Z/hpib_send_c.3i	/usr/man/cat3.Z/memcmp.3c
/usr/man/cat3.Z/hpib_spoll.3i	/usr/man/cat3.Z/memcpy.3c
/usr/man/cat3.Z/hpib_status.3i	/usr/man/cat3.Z/memory.3c
/usr/man/cat3.Z/hpib_wait_o.3i	/usr/man/cat3.Z/memset.3c
/usr/man/cat3.Z/hsearch.3c	/usr/man/cat3.Z/mfirt0.o.3
/usr/man/cat3.Z/hypot.3m	/usr/man/cat3.Z/mkfifo.3c
/usr/man/cat3.Z/idtholang.3c	/usr/man/cat3.Z/mktemp.3c
/usr/man/cat3.Z/initgroups.3c	/usr/man/cat3.Z/modify.3c
/usr/man/cat3.Z/intro.3	/usr/man/cat3.Z/monitor.3c
/usr/man/cat3.Z/io_burst.3i	/usr/man/cat3.Z/mrand48.3c
/usr/man/cat3.Z/io_eol_ctl.3i	/usr/man/cat3.Z/ndbm.3x
/usr/man/cat3.Z/nextkey.3x	/usr/man/cat3.Z/pow.3m
/usr/man/cat3.Z/nl_asctime.3c	/usr/man/cat3.Z/printf.3s



/usr/man/cat3.Z/nl_ascxtime.3c	/usr/man/cat3.Z/printmsg.3c
/usr/man/cat3.Z/nl_atof.3c	/usr/man/cat3.Z/putc.3s
/usr/man/cat3.Z/nl_conv.3c	/usr/man/cat3.Z/putchar.3s
/usr/man/cat3.Z/nl_ctime.3c	/usr/man/cat3.Z/putenv.3c
/usr/man/cat3.Z/nl_ctype.3c	/usr/man/cat3.Z/putpwent.3c
/usr/man/cat3.Z/nl_cxtime.3c	/usr/man/cat3.Z/puts.3s
/usr/man/cat3.Z/nl_fprintf.3s	/usr/man/cat3.Z/putspwent.3c
/usr/man/cat3.Z/nl_fscanf.3s	/usr/man/cat3.Z/pututline.3c
/usr/man/cat3.Z/nl_gcvt.3c	/usr/man/cat3.Z/putw.3s
/usr/man/cat3.Z/nl_init.3c	/usr/man/cat3.Z/qsort.3c
/usr/man/cat3.Z/nl_isalnum.3c	/usr/man/cat3.Z/rand.3c
/usr/man/cat3.Z/nl_isalpha.3c	/usr/man/cat3.Z/readdir.3c
/usr/man/cat3.Z/nl_iscntrl.3c	/usr/man/cat3.Z/realloc.3c
/usr/man/cat3.Z/nl_isdigit.3c	/usr/man/cat3.Z/realloc.3x
/usr/man/cat3.Z/nl_isgraph.3c	/usr/man/cat3.Z/regcmp.3x
/usr/man/cat3.Z/nl_islower.3c	/usr/man/cat3.Z/regex.3x
/usr/man/cat3.Z/nl_isprint.3c	/usr/man/cat3.Z/regexp.3x
/usr/man/cat3.Z/nl_ispunct.3c	/usr/man/cat3.Z/rewind.3s
/usr/man/cat3.Z/nl_isspace.3c	/usr/man/cat3.Z/rewinddir.3c
/usr/man/cat3.Z/nl_isupper.3c	/usr/man/cat3.Z/scanf.3s
/usr/man/cat3.Z/nl_isxdigit.3c	/usr/man/cat3.Z/secof2.3c
/usr/man/cat3.Z/nl_langinfo.3c	/usr/man/cat3.Z/seed48.3c
/usr/man/cat3.Z/nl_printf.3s	/usr/man/cat3.Z/seekdir.3c
/usr/man/cat3.Z/nl_scanf.3s	/usr/man/cat3.Z/setaclent.3c
/usr/man/cat3.Z/nl_sprintf.3s	/usr/man/cat3.Z/setaclentry.3c
/usr/man/cat3.Z/nl_sscanf.3s	/usr/man/cat3.Z/setbuf.3s
/usr/man/cat3.Z/nl_strcmp.3c	/usr/man/cat3.Z/setccent.3c
/usr/man/cat3.Z/nl_string.3c	/usr/man/cat3.Z/setfsent.3x
/usr/man/cat3.Z/nl_strncmp.3c	/usr/man/cat3.Z/setgrent.3c
/usr/man/cat3.Z/nl_strtod.3c	/usr/man/cat3.Z/setjmp.3c
/usr/man/cat3.Z/nl_tolower.3c	/usr/man/cat3.Z/setkey.3c
/usr/man/cat3.Z/nl_tools_16.3c	/usr/man/cat3.Z/setmntent.3x
/usr/man/cat3.Z/nl_toupper.3c	/usr/man/cat3.Z/setpwent.3c
/usr/man/cat3.Z/nlappend.3x	/usr/man/cat3.Z/setutent.3c
/usr/man/cat3.Z/nlcollate.3x	/usr/man/cat3.Z/setvbuf.3s
/usr/man/cat3.Z/nlconvclock.3x	/usr/man/cat3.Z/sgetl.3x
/usr/man/cat3.Z/nlconvcustd.3x	/usr/man/cat3.Z/siglongjmp.3c
/usr/man/cat3.Z/nlconvnum.3x	/usr/man/cat3.Z/sigam.3m
/usr/man/cat3.Z/nlfindstr.3x	/usr/man/cat3.Z/sigsetjmp.3c
/usr/man/cat3.Z/nlfmtcal.3x	/usr/man/cat3.Z/sin.3m
/usr/man/cat3.Z/nlfmtclock.3x	/usr/man/cat3.Z/sinh.3m
/usr/man/cat3.Z/nlfmtcustda.3x	/usr/man/cat3.Z/sleep.3c

/usr/man/cat3.Z/nlfmtdate.3x	/usr/man/cat3.Z/sprintf.3s
/usr/man/cat3.Z/nlfmtlongca.3x	/usr/man/cat3.Z/sprintmsg.3c
/usr/man/cat3.Z/nlfmtnum.3x	/usr/man/cat3.Z/sputl.3x
/usr/man/cat3.Z/nlgetlang.3x	/usr/man/cat3.Z/sqrt.3m
/usr/man/cat3.Z/nlinfo.3x	/usr/man/cat3.Z/srand.3c
/usr/man/cat3.Z/nlist.3c	/usr/man/cat3.Z/srand48.3c
/usr/man/cat3.Z/nljudge.3x	/usr/man/cat3.Z/sscanf.3s
/usr/man/cat3.Z/nlkeycompar.3x	/usr/man/cat3.Z/ssignal.3c
/usr/man/cat3.Z/nlnumspecc.3x	/usr/man/cat3.Z/stdio.3s
/usr/man/cat3.Z/nlrepchar.3x	/usr/man/cat3.Z/stdipc.3c
/usr/man/cat3.Z/nlscanmove.3x	/usr/man/cat3.Z/step.3x
/usr/man/cat3.Z/nlsubstr.3x	/usr/man/cat3.Z/store.3x
/usr/man/cat3.Z/nlswitchbuf.3x	/usr/man/cat3.Z/strcat.3c
/usr/man/cat3.Z/nltranslate.3x	/usr/man/cat3.Z/strchr.3c
/usr/man/cat3.Z/nrand48.3c	/usr/man/cat3.Z/strcmp.3c
/usr/man/cat3.Z/ opendir.3c	/usr/man/cat3.Z/strcmp16.3c
/usr/man/cat3.Z/optarg.3c	/usr/man/cat3.Z/strcmp8.3c
/usr/man/cat3.Z/opterr.3c	/usr/man/cat3.Z/strcpy.3c
/usr/man/cat3.Z/optind.3c	/usr/man/cat3.Z/strcspn.3c
/usr/man/cat3.Z/pclose.3s	/usr/man/cat3.Z/strdup.3c
/usr/man/cat3.Z/perror.3c	/usr/man/cat3.Z/sterror.3c
/usr/man/cat3.Z/popen.3s	/usr/man/cat3.Z/string.3c
/usr/man/cat3.Z/strlen.3c	/usr/man/cat4.Z/cpio.4
/usr/man/cat3.Z/strncat.3c	/usr/man/cat4.Z/d_passwd.4
/usr/man/cat3.Z/strncmp.3c	/usr/man/cat4.Z/dialups.4
/usr/man/cat3.Z/strncmp16.3c	/usr/man/cat4.Z/dir.4
/usr/man/cat3.Z/strncmp8.3c	/usr/man/cat4.Z/disktab.4
/usr/man/cat3.Z/strncpy.3c	/usr/man/cat4.Z/dosif.4
/usr/man/cat3.Z/strord.3c	/usr/man/cat4.Z/fs.4
/usr/man/cat3.Z/strpbrk.3c	/usr/man/cat4.Z/fspec.4
/usr/man/cat3.Z/strrchr.3c	/usr/man/cat4.Z/gettydefs.4
/usr/man/cat3.Z/strspn.3c	/usr/man/cat4.Z/group.4
/usr/man/cat3.Z/strtoacl.3c	/usr/man/cat4.Z/inittab.4
/usr/man/cat3.Z/strtoaclpat.3c	/usr/man/cat4.Z/inode.4
/usr/man/cat3.Z/strtoaclpatt.3c	/usr/man/cat4.Z/intro.4
/usr/man/cat3.Z/strtod.3c	/usr/man/cat4.Z/issue.4
/usr/man/cat3.Z/strtok.3c	/usr/man/cat4.Z/lif.4
/usr/man/cat3.Z/strtol.3c	/usr/man/cat4.Z/magic.4
/usr/man/cat3.Z/swab.3c	/usr/man/cat4.Z/master.4
/usr/man/cat3.Z/sys_errlist.3c	/usr/man/cat4.Z/mknod.4
/usr/man/cat3.Z/sys_nerr.3c	/usr/man/cat4.Z/mnttab.4
/usr/man/cat3.Z/system.3s	/usr/man/cat4.Z/model.4

/usr/man/cat3.Z/tan.3m	/usr/man/cat4.Z/nlist.4
/usr/man/cat3.Z/tanh.3m	/usr/man/cat4.Z/passwd.4
/usr/man/cat3.Z/tdelete.3c	/usr/man/cat4.Z/privgrp.4
/usr/man/cat3.Z/tellmdir.3c	/usr/man/cat4.Z/profile.4
/usr/man/cat3.Z/tempnam.3s	/usr/man/cat4.Z/ranlib.4
/usr/man/cat3.Z/termcap.3x	/usr/man/cat4.Z/rcsfile.4
/usr/man/cat3.Z/tfind.3c	/usr/man/cat4.Z/sccsfile.4
/usr/man/cat3.Z/tgetent.3x	/usr/man/cat4.Z/sdf.4
/usr/man/cat3.Z/tgetflag.3x	/usr/man/cat4.Z/symlink.4
/usr/man/cat3.Z/tgetnum.3x	/usr/man/cat4.Z/term.4
/usr/man/cat3.Z/tgetstr.3x	/usr/man/cat4.Z/terminfo.4
/usr/man/cat3.Z/tgoto.3x	/usr/man/cat4.Z/ttytype.4
/usr/man/cat3.Z/timezone.3c	/usr/man/cat4.Z/tztab.4
/usr/man/cat3.Z/tmpfile.3s	/usr/man/cat4.Z/utmp.4
/usr/man/cat3.Z/tmpnam.3s	/usr/man/cat4.Z/wtmp.4
/usr/man/cat3.Z/toascii.3c	/usr/man/cat5.Z
/usr/man/cat3.Z/tolower.3c	/usr/man/cat5.Z/acl.5
/usr/man/cat3.Z/toupper.3c	/usr/man/cat5.Z/ascii.5
/usr/man/cat3.Z/tputs.3x	/usr/man/cat5.Z/context.5
/usr/man/cat3.Z/trig.3m	/usr/man/cat5.Z/dirent.5
/usr/man/cat3.Z/tsearch.3c	/usr/man/cat5.Z/envIRON.5
/usr/man/cat3.Z/ttynAME.3c	/usr/man/cat5.Z/fcntl.5
/usr/man/cat3.Z/ttyslot.3c	/usr/man/cat5.Z/hier.5
/usr/man/cat3.Z/twalk.3c	/usr/man/cat5.Z/hpnlS.5
/usr/man/cat3.Z/tzname.3c	/usr/man/cat5.Z/intro.5
/usr/man/cat3.Z/tzset.3c	/usr/man/cat5.Z/ioctl.5
/usr/man/cat3.Z/undial.3c	/usr/man/cat5.Z/kana8.5
/usr/man/cat3.Z/ungetc.3s	/usr/man/cat5.Z/langid.5
/usr/man/cat3.Z/utmpname.3c	/usr/man/cat5.Z/man.5
/usr/man/cat3.Z/vfprintf.3s	/usr/man/cat5.Z/math.5
/usr/man/cat3.Z/vprintf.3s	/usr/man/cat5.Z/mm.5
/usr/man/cat3.Z/vsprintf.3s	/usr/man/cat5.Z/ndir.5
/usr/man/cat3.Z/y0.3m	/usr/man/cat5.Z/portnls.5
/usr/man/cat3.Z/y1.3m	/usr/man/cat5.Z/rcsintro.5
/usr/man/cat3.Z/yn.3m	/usr/man/cat5.Z/roman8.5
/usr/man/cat4.Z/a.out.4	/usr/man/cat5.Z/stat.5
/usr/man/cat4.Z/acct.4	/usr/man/cat5.Z/suffix.5
/usr/man/cat4.Z/ar.4	/usr/man/cat5.Z/term.5
/usr/man/cat4.Z/audit.4	/usr/man/cat5.Z/types.5
/usr/man/cat4.Z/bif.4	/usr/man/cat5.Z/values.5
/usr/man/cat4.Z/btmp.4	/usr/man/cat5.Z/varargs.5
/usr/man/cat4.Z/cdf.4	/usr/man/cat7.Z/console.7

/usr/man/cat4.Z/checklist.4	/usr/man/cat7.Z/ct.7
/usr/man/cat4.Z/clusterconf.4	/usr/man/cat7.Z/disk.7
/usr/man/cat4.Z/collate8.4	/usr/man/cat7.Z/graphics.7
/usr/man/cat4.Z/core.4	/usr/man/cat7.Z/hil.7
/usr/man/cat7.Z/intro.7	/usr/man/man1.Z/chfn.1
/usr/man/cat7.Z/iomap.7	/usr/man/man1.Z/chgrp.1
/usr/man/cat7.Z/kmem.7	/usr/man/man1.Z/chmod.1
/usr/man/cat7.Z/lp.7	/usr/man/man1.Z/chown.1
/usr/man/cat7.Z/mem.7	/usr/man/man1.Z/chsh.1
/usr/man/cat7.Z/modem.7	/usr/man/man1.Z/ci.1
/usr/man/cat7.Z/mt.7	/usr/man/man1.Z/clear.1
/usr/man/cat7.Z/nlio.7	/usr/man/man1.Z/cmp.1
/usr/man/cat7.Z/null.7	/usr/man/man1.Z/cnodes.1
/usr/man/cat7.Z/pty.7	/usr/man/man1.Z/co.1
/usr/man/cat7.Z/sttyV6.7	/usr/man/man1.Z/col.1
/usr/man/cat7.Z/termio.7	/usr/man/man1.Z/comb.1
/usr/man/cat7.Z/tty.7	/usr/man/man1.Z/comm.1
/usr/man/man1.Z/ada.1	/usr/man/man1.Z/compact.1
/usr/man/man1.Z/ada.fmgr.1	/usr/man/man1.Z/compress.1
/usr/man/man1.Z/ada.format.1	/usr/man/man1.Z/cp.1
/usr/man/man1.Z/ada.mklib.1	/usr/man/man1.Z/cpio.1
/usr/man/man1.Z/ada.mvlib.1	/usr/man/man1.Z/cpp.1
/usr/man/man1.Z/ada.probe.1	/usr/man/man1.Z/cps.1
/usr/man/man1.Z/ada.protect.1	/usr/man/man1.Z/crontab.1
/usr/man/man1.Z/ada.rmlib.1	/usr/man/man1.Z/crypt.1
/usr/man/man1.Z/ada.umgr.1	/usr/man/man1.Z/csh.1
/usr/man/man1.Z/ada.unlock.1	/usr/man/man1.Z/csplitt.1
/usr/man/man1.Z/ada.xref.1	/usr/man/man1.Z/ct.1
/usr/man/man1.Z/adb.1	/usr/man/man1.Z/ctags.1
/usr/man/man1.Z/adjust.1	/usr/man/man1.Z/cu.1
/usr/man/man1.Z/admin.1	/usr/man/man1.Z/cut.1
/usr/man/man1.Z/ar.1	/usr/man/man1.Z/cxref.1
/usr/man/man1.Z/as.1	/usr/man/man1.Z/date.1
/usr/man/man1.Z/asa.1	/usr/man/man1.Z/dc.1
/usr/man/man1.Z/astrn.1	/usr/man/man1.Z/dd.1
/usr/man/man1.Z/at.1	/usr/man/man1.Z/delta.1
/usr/man/man1.Z/atime.1	/usr/man/man1.Z/deroff.1
/usr/man/man1.Z/atrans.1	/usr/man/man1.Z/diff.1
/usr/man/man1.Z/awk.1	/usr/man/man1.Z/diff3.1
/usr/man/man1.Z/banner.1	/usr/man/man1.Z/diffh.1
/usr/man/man1.Z/basename.1	/usr/man/man1.Z/diffmk.1
/usr/man/man1.Z/basic.1	/usr/man/man1.Z/dircmp.1

/usr/man/man1.Z/batch.1	/usr/man/man1.Z/dirname.1
/usr/man/man1.Z/bc.1	/usr/man/man1.Z/disable.1
/usr/man/man1.Z/bdiff.1	/usr/man/man1.Z/dos2ux.1
/usr/man/man1.Z/bfs.1	/usr/man/man1.Z/doschmod.1
/usr/man/man1.Z/bifchgrp.1	/usr/man/man1.Z/doscp.1
/usr/man/man1.Z/bifchmod.1	/usr/man/man1.Z/dosdf.1
/usr/man/man1.Z/bifchown.1	/usr/man/man1.Z/dosll.1
/usr/man/man1.Z/bifcp.1	/usr/man/man1.Z/dosls.1
/usr/man/man1.Z/biffind.1	/usr/man/man1.Z/dosmkdir.1
/usr/man/man1.Z/bifls.1	/usr/man/man1.Z/dosrm.1
/usr/man/man1.Z/bifmkdir.1	/usr/man/man1.Z/dosrmdir.1
/usr/man/man1.Z/bifrm.1	/usr/man/man1.Z/du.1
/usr/man/man1.Z/bifrmkdir.1	/usr/man/man1.Z/dumpmsg.1
/usr/man/man1.Z/bs.1	/usr/man/man1.Z/echo.1
/usr/man/man1.Z/cal.1	/usr/man/man1.Z/ed.1
/usr/man/man1.Z/calendar.1	/usr/man/man1.Z/edit.1
/usr/man/man1.Z/cancel.1	/usr/man/man1.Z/egrep.1
/usr/man/man1.Z/cat.1	/usr/man/man1.Z/enable.1
/usr/man/man1.Z/cb.1	/usr/man/man1.Z/env.1
/usr/man/man1.Z/cc.1	/usr/man/man1.Z/ex.1
/usr/man/man1.Z/ccat.1	/usr/man/man1.Z/expand.1
/usr/man/man1.Z/cd.1	/usr/man/man1.Z/expr.1
/usr/man/man1.Z/cdb.1	/usr/man/man1.Z/f77.1
/usr/man/man1.Z/cdc.1	/usr/man/man1.Z/factor.1
/usr/man/man1.Z/cflow.1	/usr/man/man1.Z/false.1
/usr/man/man1.Z/chacl.1	/usr/man/man1.Z/fc.1
/usr/man/man1.Z/chatr.1	/usr/man/man1.Z/fdb.1
/usr/man/man1.Z/checknr.1	/usr/man/man1.Z/fgrep.1
/usr/man/man1.Z/file.1	/usr/man/man1.Z/lpstat.1
/usr/man/man1.Z/find.1	/usr/man/man1.Z/ls.1
/usr/man/man1.Z/findmsg.1	/usr/man/man1.Z/lsacl.1
/usr/man/man1.Z/findstr.1	/usr/man/man1.Z/lsf.1
/usr/man/man1.Z/finger.1	/usr/man/man1.Z/lsr.1
/usr/man/man1.Z/fixman.1	/usr/man/man1.Z/lxs.1
/usr/man/man1.Z/flint.1	/usr/man/man1.Z/m4.1
/usr/man/man1.Z/fold.1	/usr/man/man1.Z/machid.1
/usr/man/man1.Z/forder.1	/usr/man/man1.Z/mail.1
/usr/man/man1.Z/from.1	/usr/man/man1.Z/mailx.1
/usr/man/man1.Z/fsplit.1	/usr/man/man1.Z/make.1
/usr/man/man1.Z/ftio.1	/usr/man/man1.Z/makekey.1
/usr/man/man1.Z/gencat.1	/usr/man/man1.Z/man.1
/usr/man/man1.Z/get.1	/usr/man/man1.Z/mediainit.1

/usr/man/man1.Z/getaccess.1	/usr/man/man1.Z/merge.1
/usr/man/man1.Z/getcontext.1	/usr/man/man1.Z/mesg.1
/usr/man/man1.Z/getopt.1	/usr/man/man1.Z/mkdir.1
/usr/man/man1.Z/getprivgrp.1	/usr/man/man1.Z/mkmf.1
/usr/man/man1.Z/gprof.1	/usr/man/man1.Z/mkstr.1
/usr/man/man1.Z/grep.1	/usr/man/man1.Z/mkuupath.1
/usr/man/man1.Z/grget.1	/usr/man/man1.Z/mm.1
/usr/man/man1.Z/groups.1	/usr/man/man1.Z/more.1
/usr/man/man1.Z/hashcheck.1	/usr/man/man1.Z/mt.1
/usr/man/man1.Z/hashmake.1	/usr/man/man1.Z/mv.1
/usr/man/man1.Z/head.1	/usr/man/man1.Z/neqn.1
/usr/man/man1.Z/help.1	/usr/man/man1.Z/newform.1
/usr/man/man1.Z/hostname.1	/usr/man/man1.Z/newgrp.1
/usr/man/man1.Z/hp.1	/usr/man/man1.Z/news.1
/usr/man/man1.Z/hp9000s200.1	/usr/man/man1.Z/nice.1
/usr/man/man1.Z/hp9000s300.1	/usr/man/man1.Z/nl.1
/usr/man/man1.Z/hp9000s500.1	/usr/man/man1.Z/nlio.1
/usr/man/man1.Z/hp9000s800.1	/usr/man/man1.Z/nlioenv.1
/usr/man/man1.Z/hyphen.1	/usr/man/man1.Z/nliostart.1
/usr/man/man1.Z/id.1	/usr/man/man1.Z/nljust.1
/usr/man/man1.Z/ident.1	/usr/man/man1.Z/nm.1
/usr/man/man1.Z/insertmsg.1	/usr/man/man1.Z/nohup.1
/usr/man/man1.Z/intro.1	/usr/man/man1.Z/nroff.1
/usr/man/man1.Z/inv.1	/usr/man/man1.Z/od.1
/usr/man/man1.Z/ipcrm.1	/usr/man/man1.Z/pack.1
/usr/man/man1.Z/ipcs.1	/usr/man/man1.Z/page.1
/usr/man/man1.Z/isql.1	/usr/man/man1.Z/pam.1
/usr/man/man1.Z/join.1	/usr/man/man1.Z/passwd.1
/usr/man/man1.Z/kermit.1	/usr/man/man1.Z/paste.1
/usr/man/man1.Z/kill.1	/usr/man/man1.Z/pathalias.1
/usr/man/man1.Z/ksh.1	/usr/man/man1.Z/pc.1
/usr/man/man1.Z/l.1	/usr/man/man1.Z/pcat.1
/usr/man/man1.Z/last.1	/usr/man/man1.Z/pdb.1
/usr/man/man1.Z/lastb.1	/usr/man/man1.Z/pdp11.1
/usr/man/man1.Z/lastcomm.1	/usr/man/man1.Z/pg.1
/usr/man/man1.Z/ld.1	/usr/man/man1.Z/pr.1
/usr/man/man1.Z/leave.1	/usr/man/man1.Z/prealloc.1
/usr/man/man1.Z/lex.1	/usr/man/man1.Z/primes.1
/usr/man/man1.Z/lifcp.1	/usr/man/man1.Z/printenv.1
/usr/man/man1.Z/lifinit.1	/usr/man/man1.Z/prmail.1
/usr/man/man1.Z/lifls.1	/usr/man/man1.Z/prof.1
/usr/man/man1.Z/lifrename.1	/usr/man/man1.Z/prs.1

```

/usr/man/man1.Z/lifrm.1
/usr/man/man1.Z/line.1
/usr/man/man1.Z/lint.1
/usr/man/man1.Z/ll.1
/usr/man/man1.Z/ln.1
/usr/man/man1.Z/lock.1
/usr/man/man1.Z/login.1
/usr/man/man1.Z/logname.1
/usr/man/man1.Z/lorder.1
/usr/man/man1.Z/lp.1
/usr/man/man1.Z/rcs.1
/usr/man/man1.Z/rcsdiff.1
/usr/man/man1.Z/rcsmmerge.1
/usr/man/man1.Z/red.1
/usr/man/man1.Z/rev.1
/usr/man/man1.Z/rksh.1
/usr/man/man1.Z/rlog.1
/usr/man/man1.Z/rm.1
/usr/man/man1.Z/rmail.1
/usr/man/man1.Z/rmdel.1
/usr/man/man1.Z/rmdir.1
/usr/man/man1.Z/rmnl.1
/usr/man/man1.Z/rsh.1
/usr/man/man1.Z/rtprio.1
/usr/man/man1.Z/sact.1
/usr/man/man1.Z/sccsdiff.1
/usr/man/man1.Z/script.1
/usr/man/man1.Z/sdfchgrp.1
/usr/man/man1.Z/sdfchmod.1
/usr/man/man1.Z/sdfchown.1
/usr/man/man1.Z/sdfcp.1
/usr/man/man1.Z/sdffind.1
/usr/man/man1.Z/sdfll.1
/usr/man/man1.Z/sdfln.1
/usr/man/man1.Z/sdfls.1
/usr/man/man1.Z/sdfmkdir.1
/usr/man/man1.Z/sdfmv.1
/usr/man/man1.Z/sdfmrm.1
/usr/man/man1.Z/sdfmrmrmdir.1
/usr/man/man1.Z/sdiff.1
/usr/man/man1.Z/sed.1
/usr/man/man1.Z/sh.1
/usr/man/man1.Z/ps.1
/usr/man/man1.Z/psqlc.1
/usr/man/man1.Z/psqlcbl.1
/usr/man/man1.Z/psqlfor.1
/usr/man/man1.Z/psqlpas.1
/usr/man/man1.Z/ptx.1
/usr/man/man1.Z/pty.1
/usr/man/man1.Z/pwd.1
/usr/man/man1.Z/pwget.1
/usr/man/man1.Z/ratfor.1
/usr/man/man1.Z/u3b5.1
/usr/man/man1.Z/ul.1
/usr/man/man1.Z/umask.1
/usr/man/man1.Z/umodem.1
/usr/man/man1.Z/uname.1
/usr/man/man1.Z/uncompact.1
/usr/man/man1.Z/uncompress.1
/usr/man/man1.Z/unexpand.1
/usr/man/man1.Z/unget.1
/usr/man/man1.Z/uniq.1
/usr/man/man1.Z/units.1
/usr/man/man1.Z/unpack.1
/usr/man/man1.Z/uptime.1
/usr/man/man1.Z/users.1
/usr/man/man1.Z/uucp.1
/usr/man/man1.Z/uulog.1
/usr/man/man1.Z/uuname.1
/usr/man/man1.Z/uupath.1
/usr/man/man1.Z/uupick.1
/usr/man/man1.Z/uustat.1
/usr/man/man1.Z/uuto.1
/usr/man/man1.Z/uux.1
/usr/man/man1.Z/ux2dos.1
/usr/man/man1.Z/val.1
/usr/man/man1.Z/vax.1
/usr/man/man1.Z/vc.1
/usr/man/man1.Z/vi.1
/usr/man/man1.Z/vis.1
/usr/man/man1.Z/vt.1
/usr/man/man1.Z/wait.1
/usr/man/man1.Z/wc.1
/usr/man/man1.Z/wdedit.1

```

```

/usr/man/man1.Z/shar.1
/usr/man/man1.Z/shl.1
/usr/man/man1.Z/showcdf.1
/usr/man/man1.Z/size.1
/usr/man/man1.Z/sleep.1
/usr/man/man1.Z/slp.1
/usr/man/man1.Z/soelim.1
/usr/man/man1.Z/sort.1
/usr/man/man1.Z/spell.1
/usr/man/man1.Z/spellin.1
/usr/man/man1.Z/split.1
/usr/man/man1.Z/sqlutil.1
/usr/man/man1.Z/ssp.1
/usr/man/man1.Z/strings.1
/usr/man/man1.Z/strips.1
/usr/man/man1.Z/stty.1
/usr/man/man1.Z/su.1
/usr/man/man1.Z/sum.1
/usr/man/man1.Z/tabs.1
/usr/man/man1.Z/tail.1
/usr/man/man1.Z/tar.1
/usr/man/man1.Z/tbl.1
/usr/man/man1.Z/tcio.1
/usr/man/man1.Z/tee.1
/usr/man/man1.Z/test.1
/usr/man/man1.Z/time.1
/usr/man/man1.Z/touch.1
/usr/man/man1.Z/tput.1
/usr/man/man1.Z/tr.1
/usr/man/man1.Z/true.1
/usr/man/man1.Z/tset.1
/usr/man/man1.Z/tsort.1
/usr/man/man1.Z/tty.1
/usr/man/man1.Z/u3b.1
/usr/man/man1m.Z/backup.1m
/usr/man/man1m.Z/bcheckrc.1m
/usr/man/man1m.Z/bdf.1m
/usr/man/man1m.Z/bifdf.1m
/usr/man/man1m.Z/biffsc.1m
/usr/man/man1m.Z/biffbdb.1m
/usr/man/man1m.Z/bifmkfs.1m
/usr/man/man1m.Z/brc.1m
/usr/man/man1.Z/wdutil.1
/usr/man/man1.Z/what.1
/usr/man/man1.Z/whereis.1
/usr/man/man1.Z/which.1
/usr/man/man1.Z/who.1
/usr/man/man1.Z/whoami.1
/usr/man/man1.Z/write.1
/usr/man/man1.Z/xargs.1
/usr/man/man1.Z/xd.1
/usr/man/man1.Z/xd.1
/usr/man/man1.Z/yacc.1
/usr/man/man1.Z/yes.1
/usr/man/man1.Z/zcat.1
/usr/man/man1m.Z/accept.1m
/usr/man/man1m.Z/acct.1m
/usr/man/man1m.Z/acctcms.1m
/usr/man/man1m.Z/acctcom.1m
/usr/man/man1m.Z/acctcon.1m
/usr/man/man1m.Z/acctcon1.1m
/usr/man/man1m.Z/acctcon2.1m
/usr/man/man1m.Z/acctdisk.1m
/usr/man/man1m.Z/acctdusg.1m
/usr/man/man1m.Z/acctmerg.1m
/usr/man/man1m.Z/accton.1m
/usr/man/man1m.Z/acctprc.1m
/usr/man/man1m.Z/acctprc1.1m
/usr/man/man1m.Z/acctprc2.1m
/usr/man/man1m.Z/acctsh.1m
/usr/man/man1m.Z/acctwtmp.1m
/usr/man/man1m.Z/audevent.1m
/usr/man/man1m.Z/audisp.1m
/usr/man/man1m.Z/audomon.1m
/usr/man/man1m.Z/audsys.1m
/usr/man/man1m.Z/audusr.1m
/usr/man/man1m.Z/nulladm.1m
/usr/man/man1m.Z/opx25.1m
/usr/man/man1m.Z/powerfail.1m
/usr/man/man1m.Z/prctmp.1m
/usr/man/man1m.Z/prdaily.1m
/usr/man/man1m.Z/prtacct.1m
/usr/man/man1m.Z/pwck.1m
/usr/man/man1m.Z/rbootd.1m

```



/usr/man/man1m.Z/captoinfo.1m	/usr/man/man1m.Z/rc.1m
/usr/man/man1m.Z/catman.1m	/usr/man/man1m.Z/rcancel.1m
/usr/man/man1m.Z/ccck.1m	/usr/man/man1m.Z/reboot.1m
/usr/man/man1m.Z/cfuser.1m	/usr/man/man1m.Z/reconfig.1m
/usr/man/man1m.Z/chargefee.1m	/usr/man/man1m.Z/reject.1m
/usr/man/man1m.Z/chroot.1m	/usr/man/man1m.Z/restore.1m
/usr/man/man1m.Z/ckpacct.1m	/usr/man/man1m.Z/revck.1m
/usr/man/man1m.Z/clri.1m	/usr/man/man1m.Z/rlp.1m
/usr/man/man1m.Z/clrsvc.1m	/usr/man/man1m.Z/rlpdaemon.1m
/usr/man/man1m.Z/cluster.1m	/usr/man/man1m.Z/rlpstat.1m
/usr/man/man1m.Z/config.1m	/usr/man/man1m.Z/runacct.1m
/usr/man/man1m.Z/convertfs.1m	/usr/man/man1m.Z/sdfdf.1m
/usr/man/man1m.Z/cpset.1m	/usr/man/man1m.Z/sdffsck.1m
/usr/man/man1m.Z/cron.1m	/usr/man/man1m.Z/sdffsdb.1m
/usr/man/man1m.Z/csp.1m	/usr/man/man1m.Z/setmnt.1m
/usr/man/man1m.Z/cwall.1m	/usr/man/man1m.Z/setprivgrp.1m
/usr/man/man1m.Z/devnm.1m	/usr/man/man1m.Z/shutacct.1m
/usr/man/man1m.Z/df.1m	/usr/man/man1m.Z/shutdown.1m
/usr/man/man1m.Z/diskinfo.1m	/usr/man/man1m.Z/startup.1m
/usr/man/man1m.Z/diskug.1m	/usr/man/man1m.Z/swapon.1m
/usr/man/man1m.Z/dmesg.1m	/usr/man/man1m.Z/sync.1m
/usr/man/man1m.Z/dodisk.1m	/usr/man/man1m.Z/syncer.1m
/usr/man/man1m.Z/dump.1m	/usr/man/man1m.Z/sysrm.1m
/usr/man/man1m.Z/fbackup.1m	/usr/man/man1m.Z/telinit.1m
/usr/man/man1m.Z/frecover.1m	/usr/man/man1m.Z/tic.1m
/usr/man/man1m.Z/fsck.1m	/usr/man/man1m.Z/tunefs.1m
/usr/man/man1m.Z/fsclean.1m	/usr/man/man1m.Z/turnacct.1m
/usr/man/man1m.Z/fsdb.1m	/usr/man/man1m.Z/umount.1m
/usr/man/man1m.Z/fstomnt.1m	/usr/man/man1m.Z/unlink.1m
/usr/man/man1m.Z/fuser.1m	/usr/man/man1m.Z/untic.1m
/usr/man/man1m.Z/fwtmp.1m	/usr/man/man1m.Z/update.1m
/usr/man/man1m.Z/getty.1m	/usr/man/man1m.Z/uuccheck.1m
/usr/man/man1m.Z/getx25.1m	/usr/man/man1m.Z/uucico.1m
/usr/man/man1m.Z/grpck.1m	/usr/man/man1m.Z/uuclean.1m
/usr/man/man1m.Z/init.1m	/usr/man/man1m.Z/uucleanup.1m
/usr/man/man1m.Z/install.1m	/usr/man/man1m.Z/uugetty.1m
/usr/man/man1m.Z/intro.1m	/usr/man/man1m.Z/uuls.1m
/usr/man/man1m.Z/killall.1m	/usr/man/man1m.Z/uusched.1m
/usr/man/man1m.Z/lastlogin.1m	/usr/man/man1m.Z/uusnap.1m
/usr/man/man1m.Z/link.1m	/usr/man/man1m.Z/uusnaps.1m
/usr/man/man1m.Z/lpadmin.1m	/usr/man/man1m.Z/uusub.1m
/usr/man/man1m.Z/lpmove.1m	/usr/man/man1m.Z/uuxqt.1m

/usr/man/man1m.Z/lpsched.1m	/usr/man/man1m.Z/vipw.1m
/usr/man/man1m.Z/lpshut.1m	/usr/man/man1m.Z/vtdaemon.1m
/usr/man/man1m.Z/lzdev.1m	/usr/man/man1m.Z/wall.1m
/usr/man/man1m.Z/makecdf.1m	/usr/man/man1m.Z/wghodo.1m
/usr/man/man1m.Z/mkdev.1m	/usr/man/man1m.Z/wtmpfix.1m
/usr/man/man1m.Z/mkfs.1m	/usr/man/man2.Z/_exit.2
/usr/man/man1m.Z/mklost+foun.1m	/usr/man/man2.Z/access.2
/usr/man/man1m.Z/mklp.1m	/usr/man/man2.Z/acct.2
/usr/man/man1m.Z/mknod.1m	/usr/man/man2.Z/alarm.2
/usr/man/man1m.Z/mkrs.1m	/usr/man/man2.Z/audctl.2
/usr/man/man1m.Z/monacct.1m	/usr/man/man2.Z/audswitch.2
/usr/man/man1m.Z/mount.1m	/usr/man/man2.Z/audwrite.2
/usr/man/man1m.Z/mvdir.1m	/usr/man/man2.Z/brk.2
/usr/man/man1m.Z/ncheck.1m	/usr/man/man2.Z/bsdproc.2
/usr/man/man1m.Z/newfs.1m	/usr/man/man2.Z/chdir.2
/usr/man/man1m.Z/nlioinit.1m	/usr/man/man2.Z/chmod.2
/usr/man/man2.Z/chown.2	/usr/man/man2.Z/msgget.2
/usr/man/man2.Z/chroot.2	/usr/man/man2.Z/msgop.2
/usr/man/man2.Z/close.2	/usr/man/man2.Z/msgrcv.2
/usr/man/man2.Z/cnodeid.2	/usr/man/man2.Z/msgsnd.2
/usr/man/man2.Z/cnodes.2	/usr/man/man2.Z/nice.2
/usr/man/man2.Z/creat.2	/usr/man/man2.Z/open.2
/usr/man/man2.Z/dup.2	/usr/man/man2.Z/pathconf.2
/usr/man/man2.Z/dup2.2	/usr/man/man2.Z/pause.2
/usr/man/man2.Z/errno.2	/usr/man/man2.Z/pipe.2
/usr/man/man2.Z/exec.2	/usr/man/man2.Z/plock.2
/usr/man/man2.Z/execl.2	/usr/man/man2.Z/prealloc.2
/usr/man/man2.Z/execl.2	/usr/man/man2.Z/profil.2
/usr/man/man2.Z/execlp.2	/usr/man/man2.Z/ptrace.2
/usr/man/man2.Z/execv.2	/usr/man/man2.Z/read.2
/usr/man/man2.Z/execve.2	/usr/man/man2.Z/readlink.2
/usr/man/man2.Z/execvp.2	/usr/man/man2.Z/readv.2
/usr/man/man2.Z/exit.2	/usr/man/man2.Z/reboot.2
/usr/man/man2.Z/fchmod.2	/usr/man/man2.Z/rename.2
/usr/man/man2.Z/fchown.2	/usr/man/man2.Z/rmdir.2
/usr/man/man2.Z/fcntl.2	/usr/man/man2.Z/rtprio.2
/usr/man/man2.Z/fgetacl.2	/usr/man/man2.Z/sbrk.2
/usr/man/man2.Z/fgetacl.2	/usr/man/man2.Z/select.2
/usr/man/man2.Z/fork.2	/usr/man/man2.Z/semctl.2
/usr/man/man2.Z/fpathconf.2	/usr/man/man2.Z/semget.2
/usr/man/man2.Z/fsetacl.2	/usr/man/man2.Z/semop.2
/usr/man/man2.Z/fsetacl.2	/usr/man/man2.Z/setacl.2

/usr/man/man2.Z/fstat.2	/usr/man/man2.Z/setaudit.2
/usr/man/man2.Z/fstatfs.2	/usr/man/man2.Z/setaudproc.2
/usr/man/man2.Z/fsync.2	/usr/man/man2.Z/setevent.2
/usr/man/man2.Z/ftime.2	/usr/man/man2.Z/setgid.2
/usr/man/man2.Z/ftruncate.2	/usr/man/man2.Z/setgroups.2
/usr/man/man2.Z/getaccess.2	/usr/man/man2.Z/sethostname.2
/usr/man/man2.Z/getacl.2	/usr/man/man2.Z/setitimer.2
/usr/man/man2.Z/getaudit.2	/usr/man/man2.Z/setpgrp.2
/usr/man/man2.Z/getaudproc.2	/usr/man/man2.Z/setpgrp2.2
/usr/man/man2.Z/getcontext.2	/usr/man/man2.Z/setprivgrp.2
/usr/man/man2.Z/getdirentri.2	/usr/man/man2.Z/setresgid.2
/usr/man/man2.Z/getegid.2	/usr/man/man2.Z/setresuid.2
/usr/man/man2.Z/geteuid.2	/usr/man/man2.Z/settimeofda.2
/usr/man/man2.Z/getevent.2	/usr/man/man2.Z/setuid.2
/usr/man/man2.Z/getgid.2	/usr/man/man2.Z/shmat.2
/usr/man/man2.Z/getgroups.2	/usr/man/man2.Z/shmctl.2
/usr/man/man2.Z/gethostname.2	/usr/man/man2.Z/shmdt.2
/usr/man/man2.Z/getitimer.2	/usr/man/man2.Z/shmget.2
/usr/man/man2.Z/getpgrp.2	/usr/man/man2.Z/shmop.2
/usr/man/man2.Z/getpgrp2.2	/usr/man/man2.Z/sigblock.2
/usr/man/man2.Z/getpid.2	/usr/man/man2.Z/signal.2
/usr/man/man2.Z/getppid.2	/usr/man/man2.Z/sigpause.2
/usr/man/man2.Z/getprivgrp.2	/usr/man/man2.Z/sigsetmask.2
/usr/man/man2.Z/gettimeofda.2	/usr/man/man2.Z/sigspace.2
/usr/man/man2.Z/getuid.2	/usr/man/man2.Z/sigvec.2
/usr/man/man2.Z/gtty.2	/usr/man/man2.Z/sigvector.2
/usr/man/man2.Z/intro.2	/usr/man/man2.Z/stat.2
/usr/man/man2.Z/ioctl.2	/usr/man/man2.Z/statfs.2
/usr/man/man2.Z/kill.2	/usr/man/man2.Z/stime.2
/usr/man/man2.Z/killpg.2	/usr/man/man2.Z/stty.2
/usr/man/man2.Z/link.2	/usr/man/man2.Z/swapon.2
/usr/man/man2.Z/lockf.2	/usr/man/man2.Z/symlink.2
/usr/man/man2.Z/lseek.2	/usr/man/man2.Z/sync.2
/usr/man/man2.Z/lstat.2	/usr/man/man2.Z/sysconf.2
/usr/man/man2.Z/lsync.2	/usr/man/man2.Z/time.2
/usr/man/man2.Z/mkdir.2	/usr/man/man2.Z/times.2
/usr/man/man2.Z/mknod.2	/usr/man/man2.Z/truncate.2
/usr/man/man2.Z/mkrnod.2	/usr/man/man2.Z/ulimit.2
/usr/man/man2.Z/mount.2	/usr/man/man2.Z/umask.2
/usr/man/man2.Z/msgctl.2	/usr/man/man2.Z/umount.2
/usr/man/man2.Z/uname.2	/usr/man/man3.Z/crt0.3
/usr/man/man2.Z/unlink.2	/usr/man/man3.Z/crt0.o.3

/usr/man/man2.Z/ustat.2	/usr/man/man3.Z/crypt.3c
/usr/man/man2.Z/utime.2	/usr/man/man3.Z/ctermid.3s
/usr/man/man2.Z/vfork.2	/usr/man/man3.Z/ctime.3c
/usr/man/man2.Z/vfsmount.2	/usr/man/man3.Z/ctype.3c
/usr/man/man2.Z/wait.2	/usr/man/man3.Z/currlangid.3c
/usr/man/man2.Z/wait3.2	/usr/man/man3.Z/curses.3x
/usr/man/man2.Z/write.2	/usr/man/man3.Z/cuserid.3s
/usr/man/man2.Z/writev.2	/usr/man/man3.Z/cvtnum.3c
/usr/man/man3.Z/ADVANCE.3c	/usr/man/man3.Z/datalock.3c
/usr/man/man3.Z/BYTE_STATUS.3c	/usr/man/man3.Z/daylight.3c
/usr/man/man3.Z/CHARADV.3c	/usr/man/man3.Z/dbm.3x
/usr/man/man3.Z/CHARAT.3c	/usr/man/man3.Z/dbm_clearer.3x
/usr/man/man3.Z/ERROR.3x	/usr/man/man3.Z/dbm_close.3x
/usr/man/man3.Z/FIRSTof2.3c	/usr/man/man3.Z/dbm_delete.3x
/usr/man/man3.Z/GETC.3x	/usr/man/man3.Z/dbm_error.3x
/usr/man/man3.Z/INIT.3x	/usr/man/man3.Z/dbm_fetch.3x
/usr/man/man3.Z/PCHAR.3c	/usr/man/man3.Z/dbm_firstke.3x
/usr/man/man3.Z/PCHARADV.3c	/usr/man/man3.Z/dbm_nextkey.3x
/usr/man/man3.Z/PEEKC.3x	/usr/man/man3.Z/dbm_open.3x
/usr/man/man3.Z/RETURN.3x	/usr/man/man3.Z/dbm_store.3x
/usr/man/man3.Z/SECof2.3c	/usr/man/man3.Z/dbmclose.3x
/usr/man/man3.Z/UNGETC.3x	/usr/man/man3.Z/dbminit.3x
/usr/man/man3.Z/WCHAR.3c	/usr/man/man3.Z/delete.3x
/usr/man/man3.Z/WCHARADV.3c	/usr/man/man3.Z/dial.3c
/usr/man/man3.Z/_tolower.3c	/usr/man/man3.Z/directory.3c
/usr/man/man3.Z/_toupper.3c	/usr/man/man3.Z/drand48.3c
/usr/man/man3.Z/a64l.3c	/usr/man/man3.Z/ecvt.3c
/usr/man/man3.Z/abort.3c	/usr/man/man3.Z/edata.3c
/usr/man/man3.Z/abs.3c	/usr/man/man3.Z/encrypt.3c
/usr/man/man3.Z/acltostr.3c	/usr/man/man3.Z/end.3c
/usr/man/man3.Z/acos.3m	/usr/man/man3.Z/endccent.3c
/usr/man/man3.Z/addmntent.3x	/usr/man/man3.Z/endsent.3x
/usr/man/man3.Z/advance.3x	/usr/man/man3.Z/endgrent.3c
/usr/man/man3.Z/almanac.3x	/usr/man/man3.Z/endlmntent.3x
/usr/man/man3.Z/asctime.3c	/usr/man/man3.Z/endlpwent.3c
/usr/man/man3.Z/asin.3m	/usr/man/man3.Z/endutent.3c
/usr/man/man3.Z/assert.3x	/usr/man/man3.Z/erand48.3c
/usr/man/man3.Z/atan.3m	/usr/man/man3.Z/erf.3m
/usr/man/man3.Z/atan2.3m	/usr/man/man3.Z/erfc.3m
/usr/man/man3.Z/atof.3c	/usr/man/man3.Z/errno.3c
/usr/man/man3.Z/atoi.3c	/usr/man/man3.Z/etext.3c
/usr/man/man3.Z/atol.3c	/usr/man/man3.Z/exp.3m

/usr/man/man3.Z/bessel.3m	/usr/man/man3.Z/fabs.3m
/usr/man/man3.Z/bsearch.3c	/usr/man/man3.Z/fclose.3s
/usr/man/man3.Z/byte_status.3c	/usr/man/man3.Z/fcpacl.3c
/usr/man/man3.Z/calendar.3x	/usr/man/man3.Z/fcpacl.3c
/usr/man/man3.Z/calloc.3c	/usr/man/man3.Z/fcvt.3c
/usr/man/man3.Z/calloc.3x	/usr/man/man3.Z/fdopen.3s
/usr/man/man3.Z/catclose.3c	/usr/man/man3.Z/feof.3s
/usr/man/man3.Z/catgetmsg.3c	/usr/man/man3.Z/ferror.3s
/usr/man/man3.Z/catgets.3c	/usr/man/man3.Z/fetch.3x
/usr/man/man3.Z/catopen.3c	/usr/man/man3.Z/fflush.3s
/usr/man/man3.Z/catread.3c	/usr/man/man3.Z/fgetc.3s
/usr/man/man3.Z/ceil.3m	/usr/man/man3.Z/fgetcct.3c
/usr/man/man3.Z/chownacl.3c	/usr/man/man3.Z/fgetgrent.3c
/usr/man/man3.Z/clearerr.3s	/usr/man/man3.Z/fgetpwent.3c
/usr/man/man3.Z/clock.3c	/usr/man/man3.Z/fgets.3s
/usr/man/man3.Z/clock.3x	/usr/man/man3.Z/fileno.3s
/usr/man/man3.Z/closedir.3c	/usr/man/man3.Z/firstkey.3x
/usr/man/man3.Z/compile.3x	/usr/man/man3.Z/firstof2.3c
/usr/man/man3.Z/conv.3c	/usr/man/man3.Z/floor.3m
/usr/man/man3.Z/cos.3m	/usr/man/man3.Z/fmod.3m
/usr/man/man3.Z/cosh.3m	/usr/man/man3.Z/fopen.3s
/usr/man/man3.Z/cpacl.3c	/usr/man/man3.Z/fprintf.3s
/usr/man/man3.Z/fprintmsg.3c	/usr/man/man3.Z/hpib_rqst_s.3i
/usr/man/man3.Z/fputc.3s	/usr/man/man3.Z/hpib_send_c.3i
/usr/man/man3.Z/fputs.3s	/usr/man/man3.Z/hpib_spoll.3i
/usr/man/man3.Z/fread.3s	/usr/man/man3.Z/hpib_status.3i
/usr/man/man3.Z/free.3c	/usr/man/man3.Z/hpib_wait_o.3i
/usr/man/man3.Z/free.3x	/usr/man/man3.Z/hsearch.3c
/usr/man/man3.Z/freopen.3s	/usr/man/man3.Z/hypot.3m
/usr/man/man3.Z/frex.3c	/usr/man/man3.Z/idtolang.3c
/usr/man/man3.Z/frt0.o.3	/usr/man/man3.Z/initgroups.3c
/usr/man/man3.Z/fscanf.3s	/usr/man/man3.Z/intro.3
/usr/man/man3.Z/fseek.3s	/usr/man/man3.Z/io_burst.3i
/usr/man/man3.Z/fsetaclentr.3c	/usr/man/man3.Z/io_eol_ctl.3i
/usr/man/man3.Z/ftell.3s	/usr/man/man3.Z/io_get_term.3i
/usr/man/man3.Z/ftw.3c	/usr/man/man3.Z/io_interrup.3i
/usr/man/man3.Z/ftwh.3c	/usr/man/man3.Z/io_lock.3i
/usr/man/man3.Z/fwrite.3s	/usr/man/man3.Z/io_on_inter.3i
/usr/man/man3.Z/gamma.3m	/usr/man/man3.Z/io_reset.3i
/usr/man/man3.Z/gcvt.3c	/usr/man/man3.Z/io_speed_ct.3i
/usr/man/man3.Z/getc.3s	/usr/man/man3.Z/io_timeout.3i
/usr/man/man3.Z/getccid.3c	/usr/man/man3.Z/io_unlock.3i

/usr/man/man3.Z/getccent.3c	/usr/man/man3.Z/io_width_ct.3i
/usr/man/man3.Z/getccnam.3c	/usr/man/man3.Z/is_68010_pr.3c
/usr/man/man3.Z/getcdf.3c	/usr/man/man3.Z/is_68881_pr.3c
/usr/man/man3.Z/getchar.3s	/usr/man/man3.Z/is_98248A_p.3c
/usr/man/man3.Z/getcwd.3c	/usr/man/man3.Z/is_98635A_p.3c
/usr/man/man3.Z/getenv.3c	/usr/man/man3.Z/is_hw_prese.3c
/usr/man/man3.Z/getfsent.3x	/usr/man/man3.Z/isalnum.3c
/usr/man/man3.Z/getfsfile.3x	/usr/man/man3.Z/isalpha.3c
/usr/man/man3.Z/getfsspec.3x	/usr/man/man3.Z/isascii.3c
/usr/man/man3.Z/getfstype.3x	/usr/man/man3.Z/isatty.3c
/usr/man/man3.Z/getgrent.3c	/usr/man/man3.Z/iscntrl.3c
/usr/man/man3.Z/getgrgid.3c	/usr/man/man3.Z/isdigit.3c
/usr/man/man3.Z/getgrnam.3c	/usr/man/man3.Z/isgraph.3c
/usr/man/man3.Z/gethcwd.3c	/usr/man/man3.Z/islower.3c
/usr/man/man3.Z/getlogin.3c	/usr/man/man3.Z/isprint.3c
/usr/man/man3.Z/getmntent.3x	/usr/man/man3.Z/ispunct.3c
/usr/man/man3.Z/getmsg.3c	/usr/man/man3.Z/isspace.3c
/usr/man/man3.Z/getopt.3c	/usr/man/man3.Z/isupper.3c
/usr/man/man3.Z/getpass.3c	/usr/man/man3.Z/isxdigit.3c
/usr/man/man3.Z/getpw.3c	/usr/man/man3.Z/j0.3m
/usr/man/man3.Z/getpwent.3c	/usr/man/man3.Z/j1.3m
/usr/man/man3.Z/getpwnam.3c	/usr/man/man3.Z/jn.3m
/usr/man/man3.Z/getpwuid.3c	/usr/man/man3.Z/jrand48.3c
/usr/man/man3.Z/gets.3s	/usr/man/man3.Z/l3tol.3c
/usr/man/man3.Z/getspwent.3c	/usr/man/man3.Z/l64a.3c
/usr/man/man3.Z/getut.3c	/usr/man/man3.Z/langinfo.3c
/usr/man/man3.Z/getutent.3c	/usr/man/man3.Z/langinit.3c
/usr/man/man3.Z/getutid.3c	/usr/man/man3.Z/langtoid.3c
/usr/man/man3.Z/getutline.3c	/usr/man/man3.Z/lcong48.3c
/usr/man/man3.Z/getw.3s	/usr/man/man3.Z/ldexp.3c
/usr/man/man3.Z/gmtime.3c	/usr/man/man3.Z/lfind.3c
/usr/man/man3.Z/gpio_get_st.3i	/usr/man/man3.Z/localtime.3c
/usr/man/man3.Z/gpio_set_ct.3i	/usr/man/man3.Z/log.3m
/usr/man/man3.Z/gsignal.3c	/usr/man/man3.Z/log10.3m
/usr/man/man3.Z/hasmntopt.3x	/usr/man/man3.Z/logname.3c
/usr/man/man3.Z/hcreate.3c	/usr/man/man3.Z/longjmp.3c
/usr/man/man3.Z/hdestroy.3c	/usr/man/man3.Z/lrand48.3c
/usr/man/man3.Z/hpib_abort.3i	/usr/man/man3.Z/lsearch.3c
/usr/man/man3.Z/hpib_bus_st.3i	/usr/man/man3.Z/lto13.3c
/usr/man/man3.Z/hpib_card_p.3i	/usr/man/man3.Z/mallinfo.3x
/usr/man/man3.Z/hpib_eoi_ct.3i	/usr/man/man3.Z/malloc.3c
/usr/man/man3.Z/hpib_io.3i	/usr/man/man3.Z/malloc.3x

```

/usr/man/man3.Z/hpib_pass_c.3i
/usr/man/man3.Z/hpib_ppoll.3i
/usr/man/man3.Z/hpib_ppoll_.3i
/usr/man/man3.Z/hpib_ren_ct.3i
/usr/man/man3.Z/memchr.3c
/usr/man/man3.Z/memcmp.3c
/usr/man/man3.Z/memcpy.3c
/usr/man/man3.Z/memory.3c
/usr/man/man3.Z/memset.3c
/usr/man/man3.Z/mfirt0.o.3
/usr/man/man3.Z/mkfifo.3c
/usr/man/man3.Z/mktemp.3c
/usr/man/man3.Z/modf.3c
/usr/man/man3.Z/monitor.3c
/usr/man/man3.Z/mrand48.3c
/usr/man/man3.Z/ndbm.3x
/usr/man/man3.Z/nextkey.3x
/usr/man/man3.Z/nl_asctime.3c
/usr/man/man3.Z/nl_ascxtime.3c
/usr/man/man3.Z/nl_atof.3c
/usr/man/man3.Z/nl_conv.3c
/usr/man/man3.Z/nl_ctime.3c
/usr/man/man3.Z/nl_ctype.3c
/usr/man/man3.Z/nl_cxtime.3c
/usr/man/man3.Z/nl_fprintf.3s
/usr/man/man3.Z/nl_fscanf.3s
/usr/man/man3.Z/nl_gcvt.3c
/usr/man/man3.Z/nl_init.3c
/usr/man/man3.Z/nl_isalnum.3c
/usr/man/man3.Z/nl_isalpha.3c
/usr/man/man3.Z/nl_iscntrl.3c
/usr/man/man3.Z/nl_isdigit.3c
/usr/man/man3.Z/nl_isgraph.3c
/usr/man/man3.Z/nl_islower.3c
/usr/man/man3.Z/nl_isprint.3c
/usr/man/man3.Z/nl_isspace.3c
/usr/man/man3.Z/nl_isupper.3c
/usr/man/man3.Z/nl_isxdigit.3c
/usr/man/man3.Z/nl_langinfo.3c
/usr/man/man3.Z/nl_printf.3s
/usr/man/man3.Z/nl_scanf.3s
/usr/man/man3.Z/mallopt.3x
/usr/man/man3.Z/matherr.3m
/usr/man/man3.Z/mcrt0.o.3
/usr/man/man3.Z/memccpy.3c
/usr/man/man3.Z/nlscanmove.3x
/usr/man/man3.Z/nlsubstr.3x
/usr/man/man3.Z/nlswitchbuf.3x
/usr/man/man3.Z/nltranslate.3x
/usr/man/man3.Z/nrand48.3c
/usr/man/man3.Z/ opendir.3c
/usr/man/man3.Z/optarg.3c
/usr/man/man3.Z/opterr.3c
/usr/man/man3.Z/optind.3c
/usr/man/man3.Z/pclose.3s
/usr/man/man3.Z/perror.3c
/usr/man/man3.Z/popen.3s
/usr/man/man3.Z/pow.3m
/usr/man/man3.Z/printf.3s
/usr/man/man3.Z/printmsg.3c
/usr/man/man3.Z/putc.3s
/usr/man/man3.Z/putchar.3s
/usr/man/man3.Z/putenv.3c
/usr/man/man3.Z/putpwent.3c
/usr/man/man3.Z/puts.3s
/usr/man/man3.Z/putspwent.3c
/usr/man/man3.Z/pututline.3c
/usr/man/man3.Z/putw.3s
/usr/man/man3.Z/qsort.3c
/usr/man/man3.Z/rand.3c
/usr/man/man3.Z/readdir.3c
/usr/man/man3.Z/realloc.3c
/usr/man/man3.Z/realloc.3x
/usr/man/man3.Z/regcmp.3x
/usr/man/man3.Z/regex.3x
/usr/man/man3.Z/regexp.3x
/usr/man/man3.Z/rewind.3s
/usr/man/man3.Z/rewinddir.3c
/usr/man/man3.Z/scanf.3s
/usr/man/man3.Z/secof2.3c
/usr/man/man3.Z/seed48.3c
/usr/man/man3.Z/seekdir.3c
/usr/man/man3.Z/setaclent.3c

```

```

/usr/man/man3.Z/nl_sprintf.3s
/usr/man/man3.Z/nl_sscanf.3s
/usr/man/man3.Z/nl_strcmp.3c
/usr/man/man3.Z/nl_string.3c
/usr/man/man3.Z/nl_strncmp.3c
/usr/man/man3.Z/nl_strtod.3c
/usr/man/man3.Z/nl_tolower.3c
/usr/man/man3.Z/nl_tools_16.3c
/usr/man/man3.Z/nl_toupper.3c
/usr/man/man3.Z/nlappend.3x
/usr/man/man3.Z/nlcollate.3x
/usr/man/man3.Z/nlconvclock.3x
/usr/man/man3.Z/nlconvcustd.3x
/usr/man/man3.Z/nlconvnum.3x
/usr/man/man3.Z/nlfindstr.3x
/usr/man/man3.Z/nlfmtcal.3x
/usr/man/man3.Z/nlfmtclock.3x
/usr/man/man3.Z/nlfmtcustda.3x
/usr/man/man3.Z/nlfmtdate.3x
/usr/man/man3.Z/nlfmtlongca.3x
/usr/man/man3.Z/nlfmtnum.3x
/usr/man/man3.Z/nlgetlang.3x
/usr/man/man3.Z/nlinfo.3x
/usr/man/man3.Z/nlist.3c
/usr/man/man3.Z/nljudge.3x
/usr/man/man3.Z/nlkeycompar.3x
/usr/man/man3.Z/nlnumspec.3x
/usr/man/man3.Z/nlrepchar.3x
/usr/man/man3.Z/step.3x
/usr/man/man3.Z/store.3x
/usr/man/man3.Z/strcat.3c
/usr/man/man3.Z/strchr.3c
/usr/man/man3.Z/strcmp.3c
/usr/man/man3.Z/strcmp16.3c
/usr/man/man3.Z/strcmp8.3c
/usr/man/man3.Z/strcpy.3c
/usr/man/man3.Z/strcspn.3c
/usr/man/man3.Z/strdup.3c
/usr/man/man3.Z/strerror.3c
/usr/man/man3.Z/string.3c
/usr/man/man3.Z/strlen.3c
/usr/man/man3.Z/strncat.3c
/usr/man/man3.Z/setaclentry.3c
/usr/man/man3.Z/setbuf.3s
/usr/man/man3.Z/setccent.3c
/usr/man/man3.Z/setfsent.3x
/usr/man/man3.Z/setgrent.3c
/usr/man/man3.Z/setjmp.3c
/usr/man/man3.Z/setkey.3c
/usr/man/man3.Z/setmntent.3x
/usr/man/man3.Z/setpwent.3c
/usr/man/man3.Z/setutent.3c
/usr/man/man3.Z/setvbuf.3s
/usr/man/man3.Z/sgetl.3x
/usr/man/man3.Z/siglongjmp.3c
/usr/man/man3.Z/signgam.3m
/usr/man/man3.Z/sigsetjmp.3c
/usr/man/man3.Z/sin.3m
/usr/man/man3.Z/sinh.3m
/usr/man/man3.Z/sleep.3c
/usr/man/man3.Z/sprintf.3s
/usr/man/man3.Z/sprintmsg.3c
/usr/man/man3.Z/sputl.3x
/usr/man/man3.Z/sqrt.3m
/usr/man/man3.Z/srand.3c
/usr/man/man3.Z/srand48.3c
/usr/man/man3.Z/sscanf.3s
/usr/man/man3.Z/ssignal.3c
/usr/man/man3.Z/stdio.3s
/usr/man/man3.Z/stdipc.3c
/usr/man/man3.Z/yn.3m
/usr/man/man4.Z/a.out.4
/usr/man/man4.Z/acct.4
/usr/man/man4.Z/ar.4
/usr/man/man4.Z/audit.4
/usr/man/man4.Z/bif.4
/usr/man/man4.Z/btmp.4
/usr/man/man4.Z/cdf.4
/usr/man/man4.Z/ckclist.4
/usr/man/man4.Z/clusterconf.4
/usr/man/man4.Z/collate8.4
/usr/man/man4.Z/core.4
/usr/man/man4.Z/cpio.4
/usr/man/man4.Z/d_passwd.4

```



/usr/man/man3.Z/strncmp.3c	/usr/man/man4.Z/dialups.4
/usr/man/man3.Z/strncmp16.3c	/usr/man/man4.Z/dir.4
/usr/man/man3.Z/strncmp8.3c	/usr/man/man4.Z/disktab.4
/usr/man/man3.Z/strncpy.3c	/usr/man/man4.Z/dosif.4
/usr/man/man3.Z/strord.3c	/usr/man/man4.Z/fs.4
/usr/man/man3.Z/strprbk.3c	/usr/man/man4.Z/fspec.4
/usr/man/man3.Z/strrchr.3c	/usr/man/man4.Z/gettydefs.4
/usr/man/man3.Z/strspn.3c	/usr/man/man4.Z/group.4
/usr/man/man3.Z/strtoacl.3c	/usr/man/man4.Z/inittab.4
/usr/man/man3.Z/strtoaclpat.3c	/usr/man/man4.Z/inode.4
/usr/man/man3.Z/strtoaclpatt.3c	/usr/man/man4.Z/intro.4
/usr/man/man3.Z/strtod.3c	/usr/man/man4.Z/issue.4
/usr/man/man3.Z/strtok.3c	/usr/man/man4.Z/lif.4
/usr/man/man3.Z/strtol.3c	/usr/man/man4.Z/magic.4
/usr/man/man3.Z/swab.3c	/usr/man/man4.Z/master.4
/usr/man/man3.Z/sys_errlist.3c	/usr/man/man4.Z/mknod.4
/usr/man/man3.Z/sys_nerr.3c	/usr/man/man4.Z/mnttab.4
/usr/man/man3.Z/system.3s	/usr/man/man4.Z/model.4
/usr/man/man3.Z/tan.3m	/usr/man/man4.Z/nlist.4
/usr/man/man3.Z/tanh.3m	/usr/man/man4.Z/passwd.4
/usr/man/man3.Z/tdelete.3c	/usr/man/man4.Z/privgrp.4
/usr/man/man3.Z/telldir.3c	/usr/man/man4.Z/profile.4
/usr/man/man3.Z/tempnam.3s	/usr/man/man4.Z/ranlib.4
/usr/man/man3.Z/termcap.3x	/usr/man/man4.Z/rcsfile.4
/usr/man/man3.Z/tfind.3c	/usr/man/man4.Z/sccsfile.4
/usr/man/man3.Z/tgetent.3x	/usr/man/man4.Z/sdf.4
/usr/man/man3.Z/tgetflag.3x	/usr/man/man4.Z/symlink.4
/usr/man/man3.Z/tgetnum.3x	/usr/man/man4.Z/term.4
/usr/man/man3.Z/tgetstr.3x	/usr/man/man4.Z/terminfo.4
/usr/man/man3.Z/tgoto.3x	/usr/man/man4.Z/ttytype.4
/usr/man/man3.Z/timezone.3c	/usr/man/man4.Z/tztab.4
/usr/man/man3.Z/tmpfile.3s	/usr/man/man4.Z/utmp.4
/usr/man/man3.Z/tmpnam.3s	/usr/man/man4.Z/wtmp.4
/usr/man/man3.Z/toascii.3c	/usr/man/man5.Z/acl.5
/usr/man/man3.Z/tolower.3c	/usr/man/man5.Z/ascii.5
/usr/man/man3.Z/toupper.3c	/usr/man/man5.Z/context.5
/usr/man/man3.Z/tputs.3x	/usr/man/man5.Z/dirent.5
/usr/man/man3.Z/trig.3m	/usr/man/man5.Z/envIRON.5
/usr/man/man3.Z/tsearch.3c	/usr/man/man5.Z/fcntl.5
/usr/man/man3.Z/ttyname.3c	/usr/man/man5.Z/hier.5
/usr/man/man3.Z/ttyslot.3c	/usr/man/man5.Z/hpnl.5
/usr/man/man3.Z/twalk.3c	/usr/man/man5.Z/intro.5

/usr/man/man3.Z/tzname.3c	/usr/man/man5.Z/ioctl.5
/usr/man/man3.Z/tzset.3c	/usr/man/man5.Z/kana8.5
/usr/man/man3.Z/undial.3c	/usr/man/man5.Z/langid.5
/usr/man/man3.Z/ungetc.3s	/usr/man/man5.Z/man.5
/usr/man/man3.Z/utmpname.3c	/usr/man/man5.Z/math.5
/usr/man/man3.Z/vfprintf.3s	/usr/man/man5.Z/mm.5
/usr/man/man3.Z/vprintf.3s	/usr/man/man5.Z/ndir.5
/usr/man/man3.Z/vsprintf.3s	/usr/man/man5.Z/portnls.5
/usr/man/man3.Z/y0.3m	/usr/man/man5.Z/rcintro.5
/usr/man/man3.Z/y1.3m	/usr/man/man5.Z/roman8.5
/usr/man/man5.Z/stat.5	/usr/man/man7.Z/iomap.7
/usr/man/man5.Z/suffix.5	/usr/man/man7.Z/kmem.7
/usr/man/man5.Z/term.5	/usr/man/man7.Z/lp.7
/usr/man/man5.Z/types.5	/usr/man/man7.Z/mem.7
/usr/man/man5.Z/values.5	/usr/man/man7.Z/modem.7
/usr/man/man5.Z/varargs.5	/usr/man/man7.Z/mt.7
/usr/man/man7.Z/console.7	/usr/man/man7.Z/nlio.7
/usr/man/man7.Z/ct.7	/usr/man/man7.Z/null.7
/usr/man/man7.Z/disk.7	/usr/man/man7.Z/pty.7
/usr/man/man7.Z/graphics.7	/usr/man/man7.Z/sttyV6.7
/usr/man/man7.Z/hil.7	/usr/man/man7.Z/termio.7
/usr/man/man7.Z/intro.7	/usr/man/man7.Z/tty.7

## Fileset PTEXT

/bin/hyphen	/usr/lib/ms/s.eqn
/etc/catman	/usr/lib/ms/s.ref
/etc/newconfig/hlista	/usr/lib/ms/s.tbl
/etc/newconfig/hlistb	/usr/lib/ms/s.ths
/etc/newconfig/hstop	/usr/lib/ms/s.toc
/system/PTEXT/customize	/usr/lib/spell/compress
/system/PTEXT/revlist	/usr/lib/spell/hashcheck
/usr/bin/checknr	/usr/lib/spell/hashmake
/usr/bin/deroff	/usr/lib/spell/spellin
/usr/bin/diffmk	/usr/lib/spell/spellprog
/usr/bin/fixman	/usr/lib/suftab
/usr/bin/mm	/usr/lib/term/tab2631
/usr/bin/mmchek	/usr/lib/term/tab2631-c
/usr/bin/neqn	/usr/lib/term/tab2631-e
/usr/bin/nroff	/usr/lib/term/tab300
/usr/bin/pg	/usr/lib/term/tab300-12
/usr/bin/ptx	/usr/lib/term/tab300S
/usr/bin/soelim	/usr/lib/term/tab300S-12
/usr/bin/spell	/usr/lib/term/tab300s
/usr/bin/tbl	/usr/lib/term/tab300s-12

/usr/bin/ul	/usr/lib/term/tab37
/usr/lib/eign	/usr/lib/term/tab382
/usr/lib/getNAME	/usr/lib/term/tab4000A
/usr/lib/macros/an	/usr/lib/term/tab4000a
/usr/lib/macros/cmp.n.d.an	/usr/lib/term/tab450
/usr/lib/macros/cmp.n.d.m	/usr/lib/term/tab450-12
/usr/lib/macros/cmp.n.t.an	/usr/lib/term/tab832
/usr/lib/macros/cmp.n.t.m	/usr/lib/term/tabX
/usr/lib/macros/mmn	/usr/lib/term/ta1
/usr/lib/macros/ucmp.n.an	/usr/lib/term/tabklp
/usr/lib/macros/ucmp.n.m	/usr/lib/term/tablj
/usr/lib/mkwhatis	/usr/lib/term/tablp
/usr/lib/ms/README	/usr/lib/term/tabtn300
/usr/lib/ms/end.awk	/usr/lib/tmac/tmac.an
/usr/lib/ms/endnote	/usr/lib/tmac/tmac.m
/usr/lib/ms/s.acc	/usr/lib/tmac/tmac.s
/usr/lib/ms/s.cov	

## Partition SYS\_TOOLS

### Fileset PCEUTIL

/system/PCEUTIL/customize	/usr/CE.utilities/98730/davucode/shutdown.T
/system/PCEUTIL/revlist	/usr/CE.utilities/98730/davucode/spring_d.T
/usr/CE.utilities/98720/CLEAN	/usr/CE.utilities/98730/davucode/spring_v.T
/usr/CE.utilities/98720/START	/usr/CE.utilities/98730/davucode/t_addfsp0.d
/usr/CE.utilities/98720/Taf	/usr/CE.utilities/98730/davucode/t_addfsp1.d
/usr/CE.utilities/98730/help	/usr/CE.utilities/98730/helptext
/usr/CE.utilities/98730/how_many	/usr/CE.utilities/98730/section
/usr/CE.utilities/98730/what_tests	/usr/CE.utilities/CS80/exerciser
/usr/CE.utilities/Crtadjust/adjust	/usr/CE.utilities/Crtadjust/catdevname
/usr/CE.utilities/Crtadjust/help	/usr/CE.utilities/Crtadjust/helptext
/usr/CE.utilities/Floatcard/help	/usr/CE.utilities/Floatcard/helptext
/usr/CE.utilities/Sft/98248A.test	/usr/CE.utilities/Sft/dgraphics.c
/usr/CE.utilities/Sft/disctest	/usr/CE.utilities/Sft/rungrafxtest
/usr/CE.utilities/Sft/discstest2	/usr/CE.utilities/Sft/runtst
/usr/CE.utilities/Sft/fastcat.c	/usr/CE.utilities/Sft/sft
/usr/CE.utilities/Sft/float.h	/usr/CE.utilities/Sft/soptest
/usr/CE.utilities/Sft/fltaccept.c	/usr/CE.utilities/Sft/soptest.gold
/usr/CE.utilities/Sft/get4digits	/usr/CE.utilities/Sft/tablet.c
/usr/CE.utilities/Sft/get6digits	/usr/CE.utilities/Sft/tablet.o
/usr/CE.utilities/Sft/graphics.c	/usr/CE.utilities/Sft/tablettest
/usr/CE.utilities/Sft/graphicsmenu	/usr/CE.utilities/Sft/tapemenu

```

/usr/CE.utilities/Sft/graphicstest /usr/CE.utilities/Sft/tapetest
/usr/CE.utilities/Sft/help /usr/CE.utilities/Sft/termtest
/usr/CE.utilities/Sft/helptext /usr/CE.utilities/Sft/testpattern
/usr/CE.utilities/Sft/optest /usr/CE.utilities/Sft/topmenu
/usr/CE.utilities/Sft/optest.gold /usr/CE.utilities/VME/help
/usr/CE.utilities/Sft/printermenu /usr/CE.utilities/VME/vme_README
/usr/CE.utilities/Sft/printertest /usr/CE.utilities/VME/vme_test
/usr/CE.utilities/Sft/readterm /usr/CE.utilities/help
/usr/CE.utilities/Sft/readterm.c /usr/CE.utilities/helptext
/usr/CE.utilities/Sft/rs232test /usr/CE.utilities/Sft/discmenu
/usr/CE.utilities/98720/ucode/ace.1
/usr/CE.utilities/98720/ucode/ace.2
/usr/CE.utilities/98720/ucode/ace.3
/usr/CE.utilities/98720/ucode/alu.o
/usr/CE.utilities/98720/ucode/cdma.o
/usr/CE.utilities/98720/ucode/clear.o
/usr/CE.utilities/98720/ucode/crc.o
/usr/CE.utilities/98720/ucode/dat
/usr/CE.utilities/98720/ucode/dcram.o
/usr/CE.utilities/98720/ucode/dram.o
/usr/CE.utilities/98720/ucode/fbt.o
/usr/CE.utilities/98720/ucode/fconfig.o
/usr/CE.utilities/98720/ucode/fpchips.o
/usr/CE.utilities/98720/ucode/frame_dump
/usr/CE.utilities/98720/ucode/make_crtmap
/usr/CE.utilities/98720/ucode/opt.o
/usr/CE.utilities/98720/ucode/pram.o
/usr/CE.utilities/98720/ucode/refcrc
/usr/CE.utilities/98720/ucode/refcrc_D
/usr/CE.utilities/98720/ucode/refcrc_E
/usr/CE.utilities/98720/ucode/reg.o
/usr/CE.utilities/98720/ucode/rt.o
/usr/CE.utilities/98720/ucode/seq.o
/usr/CE.utilities/98720/ucode/spininit
/usr/CE.utilities/98720/ucode/spinonce
/usr/CE.utilities/98720/ucode/td_backup
/usr/CE.utilities/98720/ucode/tecode.o
/usr/CE.utilities/98720/ucode/temp.o
/usr/CE.utilities/98720/ucode/test_disable
/usr/CE.utilities/98720/ucode/time.o
/usr/CE.utilities/98720/ucode/tri.o
/usr/CE.utilities/98730/CLEAN
/usr/CE.utilities/98730/START

```

/usr/CE.utilities/98730/davucode/t\_addfsp2.d  
/usr/CE.utilities/98730/davucode/t\_addfsp3.d  
/usr/CE.utilities/98730/davucode/t\_addfsp4.d  
/usr/CE.utilities/98730/davucode/t\_addfsp5.d  
/usr/CE.utilities/98730/davucode/t\_ccreg0.d  
/usr/CE.utilities/98730/davucode/t\_cdram0.d  
/usr/CE.utilities/98730/davucode/t\_dataram0.d  
/usr/CE.utilities/98730/davucode/t\_dcram0.d  
/usr/CE.utilities/98730/davucode/t\_divfsp0.d  
/usr/CE.utilities/98730/davucode/t\_divfsp1.d  
/usr/CE.utilities/98730/davucode/t\_divfsp2.d  
/usr/CE.utilities/98730/davucode/t\_divfsp3.d  
/usr/CE.utilities/98730/davucode/t\_divfsp4.d  
/usr/CE.utilities/98730/davucode/t\_divfsp5.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp0.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp1.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp2.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp3.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp4.d  
/usr/CE.utilities/98730/davucode/t\_mulfsp5.d  
/usr/CE.utilities/98730/davucode/t\_pramaddr.d  
/usr/CE.utilities/98730/davucode/t\_qrotate0.d  
/usr/CE.utilities/98730/davucode/t\_ready1.d  
/usr/CE.utilities/98730/davucode/t\_test10.d  
/usr/CE.utilities/98730/davucode/t\_test11.d  
/usr/CE.utilities/98730/davucode/t\_test12.d  
/usr/CE.utilities/98730/davucode/t\_test13.d  
/usr/CE.utilities/98730/davucode/t\_test14.d  
/usr/CE.utilities/98730/davucode/t\_test15.d  
/usr/CE.utilities/98730/davucode/t\_test16.d  
/usr/CE.utilities/98730/davucode/t\_test17.d  
/usr/CE.utilities/98730/davucode/t\_test23.d  
/usr/CE.utilities/98730/davucode/t\_test8.d  
/usr/CE.utilities/98730/davucode/ace1.T  
/usr/CE.utilities/98730/davucode/t\_test9.d  
/usr/CE.utilities/98730/davucode/ace2.T  
/usr/CE.utilities/98730/davucode/te\_data.d  
/usr/CE.utilities/98730/davucode/ace3.T  
/usr/CE.utilities/98730/davucode/the\_demo  
/usr/CE.utilities/98730/davucode/ace4.T  
/usr/CE.utilities/98730/davucode/torus\_d.T  
/usr/CE.utilities/98730/davucode/ace\_blue.T  
/usr/CE.utilities/98730/davucode/torus\_v.T

```

/usr/CE.utilities/98730/davucode/ace_grn.T
/usr/CE.utilities/98730/davucode/vw_d.T
/usr/CE.utilities/98730/davucode/ace_hang.T
/usr/CE.utilities/98730/davucode/vw_v.T
/usr/CE.utilities/98730/davucode/ace_rect.T
/usr/CE.utilities/98730/davucode/ace_rect1.T
/usr/CE.utilities/98730/davucode/ace_rect2.T
/usr/CE.utilities/98730/davucode/ace_rect3.T
/usr/CE.utilities/98730/davucode/ace_rect4.T
/usr/CE.utilities/98730/davucode/ace_rect5.T
/usr/CE.utilities/98730/davucode/ace_red.T
/usr/CE.utilities/98730/davucode/ace_vio.T
/usr/CE.utilities/98730/davucode/ace_white.T
/usr/CE.utilities/98730/davucode/crc.d
/usr/CE.utilities/98730/davucode/debug
/usr/CE.utilities/98730/davucode/disable.T
/usr/CE.utilities/98730/davucode/dl_script.T
/usr/CE.utilities/98730/davucode/field
/usr/CE.utilities/98730/davucode/fp.d
/usr/CE.utilities/98730/davucode/fpx.d
/usr/CE.utilities/98730/davucode/glass1_d.T
/usr/CE.utilities/98730/davucode/glass1_v.T
/usr/CE.utilities/98730/davucode/gs3ddiag.msg
/usr/CE.utilities/98730/davucode/mkcrmap.T
/usr/CE.utilities/98730/davucode/refcrc.T
/usr/CE.utilities/Crtadjust/crtadj.c
/usr/CE.utilities/Crtadjust/devmenu
/usr/CE.utilities/Floatcard/float.h
/usr/CE.utilities/Floatcard/float.mk
/usr/CE.utilities/Floatcard/fltaccept.c
/usr/CE.utilities/Floatcard/fltexhaust.c

```

## Fileset PDRIVERS

```

/SYSDEBUG /usr/lib/drivers/interface.skel
/system/PDRIVERS/revlist /usr/lib/drivers/scsi.c
/usr/lib/drivers/Makefile /usr/lib/drivers/scsi_if.c
/usr/lib/drivers/Readme /usr/lib/drivers/snalink.c
/usr/lib/drivers/ciper.c /usr/lib/drivers/vme.dma.skel
/usr/lib/drivers/gpio.c /usr/lib/drivers/vme.skel.c

```

## Fileset PFILTER

```

/bin/head /usr/bin/pack
/bin/od /usr/bin/paste

```

/bin/size	/usr/bin/pcat
/bin/xd	/usr/bin/rev
/system/PFILTER/customize	/usr/bin/rmnl
/system/PFILTER/revlist	/usr/bin/sdiff
/usr/bin/adjust	/usr/bin/split
/usr/bin/ccat	/usr/bin/ssp
/usr/bin/comm	/usr/bin/strings
/usr/bin/compact	/usr/bin/tsort
/usr/bin/diff3	/usr/bin/uncompact
/usr/bin/dircmp	/usr/bin/uniq
/usr/bin/join	/usr/bin/unpack
/usr/bin/newform	/usr/lib/diff3prog
/usr/bin/nl	/usr/lib/libndir.a

### Fileset PRLP

/system/PRLP/revlist	/usr/lib/rlpstat
/usr/lib/rcancel	/usr/spool/lp/cmodel/rcmodel
/usr/lib/rlp	/usr/spool/lp/model/rmodel
/usr/lib/rlpdaemon	/usr/spool/lp/smodel/rsmodel

### Fileset PSYSCOM

/etc/getx25	/usr/lib/uucp/Uutry
/etc/newconfig/ctCDF	/usr/lib/uucp/X25/HP2334A.clr
/etc/newconfig/cuCDF	/usr/lib/uucp/X25/HP2334A.in
/etc/newconfig/mailx.rc	/usr/lib/uucp/X25/HP2334A.out
/etc/vtdaemon	/usr/lib/uucp/X25/HP2334A.outf
/etc/vtgateway	/usr/lib/uucp/X25/HP2334A.outg
/etc/vtserver	/usr/lib/uucp/X25/HP2334A.outi
/system/PSYSCOM/CDFcustomize	/usr/lib/uucp/X25/clrsvc
/system/PSYSCOM/customize	/usr/lib/uucp/X25/opx25
/system/PSYSCOM/revlist	/usr/lib/uucp/X25/ventel.out
/usr/bin/ct	/usr/lib/uucp/newconfig/Devices
/usr/bin/cu	/usr/lib/uucp/newconfig/Dialcodes
/usr/bin/from	/usr/lib/uucp/newconfig/Dialers
/usr/bin/kermit	/usr/lib/uucp/newconfig/Maxuuscheds
/usr/bin/mailx	/usr/lib/uucp/newconfig/Maxuuxqts
/usr/bin/mkuupath	/usr/lib/uucp/newconfig/Permissions
/usr/bin/pathalias	/usr/lib/uucp/newconfig/Poll
/usr/bin/prmail	/usr/lib/uucp/newconfig/Systems
/usr/bin/modem	/usr/lib/uucp/newconfig/dialit
/usr/bin/uucp	/usr/lib/uucp/newconfig/dialit.c
/usr/bin/uulog	/usr/lib/uucp/newconfig/remote.unknown
/usr/bin/uuls	/usr/lib/uucp/newconfig/uudemmon.admin

/usr/bin/uuname	/usr/lib/uucp/newconfig/uudemon.cleau
/usr/bin/uupath	/usr/lib/uucp/newconfig/uudemon.hour
/usr/bin/uupick	/usr/lib/uucp/newconfig/uudemon.poll
/usr/bin/uusnap	/usr/lib/uucp/uucheck
/usr/bin/uusnaps	/usr/lib/uucp/uucico
/usr/bin/uustat	/usr/lib/uucp/uuclean
/usr/bin/uuto	/usr/lib/uucp/uucleanup
/usr/bin/uux	/usr/lib/uucp/uugetty
/usr/bin/vt	/usr/lib/uucp/uusched
/usr/lib/mailx/mailx.help	/usr/lib/uucp/uusub
/usr/lib/mailx/mailx.help.~	/usr/lib/uucp/uuxqt
/usr/lib/mailx/rmmail	

## Partition STARBASE

### Fileset DSTARBAS

/system/DSTARBAS/customize	/usr/lib/starbase/demos/SBUTILS/sbutils.h
/system/DSTARBAS/revlist	/usr/lib/starbase/demos/beach.c
/usr/lib/starbase/demos/DATA/frame	/usr/lib/starbase/demos/cube.c
/usr/lib/starbase/demos/DATA/glow	/usr/lib/starbase/demos/example.c
/usr/lib/starbase/demos/DATA/pgmap	/usr/lib/starbase/demos/interp_cgm.c
/usr/lib/starbase/demos/logo.c	/usr/lib/starbase/demos/liberty.c
/usr/lib/starbase/demos/makefile	/usr/lib/starbase/demos/boxes.c
/usr/lib/starbase/demos/README	/usr/lib/starbase/demos/panzoom.c
/usr/lib/starbase/demos/sbutils.h	/usr/lib/starbase/demos/SBUTILS/makefile
/usr/lib/starbase/demos/shuttle.c	/usr/lib/starbase/demos/SBUTILS/sb.3d.c
/usr/lib/starbase/demos/spots.c	/usr/lib/starbase/demos/SBUTILS/sb.glob.c
/usr/lib/starbase/demos/vases.c	/usr/lib/starbase/demos/SBUTILS/sb.glob.h
/usr/lib/starbase/demos/xwing.c	/usr/lib/starbase/demos/SBUTILS/sbutils.c.h
/usr/lib/starbase/demos/DATA/cowtping.cgm	
/usr/lib/starbase/demos/DATA/planet	
/usr/lib/starbase/demos/DATA/statue	
/usr/lib/starbase/demos/DATA/xwing_data	
/usr/lib/starbase/demos/SBUTILS/README	

### Fileset MSTARBAS

/system/MSTARBAS/customize	/usr/man/man3.Z/display_ena.3g
/system/MSTARBAS/revlist	/usr/man/man3.Z/double_buff.3g
/usr/man/man1.Z/pcltrans.1	/usr/man/man3.Z/draw.3g
/usr/man/man3.Z/append_text.3g	/usr/man/man3.Z/draw2d.3g
/usr/man/man3.Z/arc.3g	/usr/man/man3.Z/draw3d.3g
/usr/man/man3.Z/await_event.3g	/usr/man/man3.Z/drawing_mod.3g



```

/usr/man/man3.Z/await_retra.3g
/usr/man/man3.Z/backface_co.3g
/usr/man/man3.Z/background_.3g
/usr/man/man3.Z/bank_switch.3g
/usr/man/man3.Z/bezier_knot.3g
/usr/man/man3.Z/bf_control.3g
/usr/man/man3.Z/bf_fill_col.3g
/usr/man/man3.Z/bf_interior.3g
/usr/man/man3.Z/bf_perimete.3g
/usr/man/man3.Z/bf_surface_.3g
/usr/man/man3.Z/bitmap_prin.3g
/usr/man/man3.Z/bitmap_to_f.3g
/usr/man/man3.Z/block_move.3g
/usr/man/man3.Z/block_read.3g
/usr/man/man3.Z/block_write.3g
/usr/man/man3.Z/buffer_mode.3g
/usr/man/man3.Z/cgm_to_star.3g
/usr/man/man3.Z/character_e.3g
/usr/man/man3.Z/character_h.3g
/usr/man/man3.Z/character_s.3g
/usr/man/man3.Z/character_w.3g
/usr/man/man3.Z/circle.3g
/usr/man/man3.Z/clear_contr.3g
/usr/man/man3.Z/clear_view_.3g
/usr/man/man3.Z/clip_depth.3g
/usr/man/man3.Z/clip_indica.3g
/usr/man/man3.Z/clip_rectan.3g
/usr/man/man3.Z/concat_matr.3g
/usr/man/man3.Z/concat_tran.3g
/usr/man/man3.Z/curve_resol.3g
/usr/man/man3.Z/dbuffer_swi.3g
/usr/man/man3.Z/dc_to_vdc.3g
/usr/man/man3.Z/dcbitmap_pr.3g
/usr/man/man3.Z/dcbitmap_to.3g
/usr/man/man3.Z/dcblock_mov.3g
/usr/man/man3.Z/dcblock_rea.3g
/usr/man/man3.Z/dcblock_wri.3g
/usr/man/man3.Z/echo_type.3g
/usr/man/man3.Z/echo_update.3g
/usr/man/man3.Z/ellipse.3g
/usr/man/man3.Z/enable_even.3g
/usr/man/man3.Z/file_print.3g
/usr/man/man3.Z/file_to_bit.3g
/usr/man/man3.Z/file_to_dcb.3g
/usr/man/man3.Z/file_to_int.3g
/usr/man/man3.Z/fill_color.3g
/usr/man/man3.Z/fill_color_.3g
/usr/man/man3.Z/fill_dither.3g
/usr/man/man3.Z/flush_buffe.3g
/usr/man/man3.Z/flush_matri.3g
/usr/man/man3.Z/gclose.3g
/usr/man/man3.Z/gerr_control.3g
/usr/man/man3.Z/gerr_defaul.3g
/usr/man/man3.Z/gerr_messag.3g
/usr/man/man3.Z/gerr_p.3g
/usr/man/man3.Z/gerr_print_.3g
/usr/man/man3.Z/gerr_proced.3g
/usr/man/man3.Z/gescape.3g
/usr/man/man3.Z/gopen.3g
/usr/man/man3.Z/hatch_orien.3g
/usr/man/man3.Z/hatch_spaci.3g
/usr/man/man3.Z/hatch_type.3g
/usr/man/man3.Z/hidden_surf.3g
/usr/man/man3.Z/highlight.3g
/usr/man/man3.Z/highlight_a.3g
/usr/man/man3.Z/highlight_c.3g
/usr/man/man3.Z/highlight_o.3g
/usr/man/man3.Z/highlight_t.3g
/usr/man/man3.Z/hit_mode.3g
/usr/man/man3.Z/initiate_re.3g
/usr/man/man3.Z/inq_pick_de.3g
/usr/man/man3.Z/inq_pick_wi.3g
/usr/man/man3.Z/inquire_cgm.3g
/usr/man/man3.Z/inquire_col.3g

```

/usr/man/man3.Z/dcharacter.3g	/usr/man/man3.Z/inquire_cur.3g
/usr/man/man3.Z/dccircle.3g	/usr/man/man3.Z/inquire_dis.3g
/usr/man/man3.Z/dcdraw.3g	/usr/man/man3.Z/inquire_fb.3g
/usr/man/man3.Z/dcecho_type.3g	/usr/man/man3.Z/inquire_fil.3g
/usr/man/man3.Z/dcecho_upda.3g	/usr/man/man3.Z/inquire_ger.3g
/usr/man/man3.Z/dcmarker_si.3g	/usr/man/man3.Z/inquire_hit.3g
/usr/man/man3.Z/dcmove.3g	/usr/man/man3.Z/inquire_id.3g
/usr/man/man3.Z/dcpartial_p.3g	/usr/man/man3.Z/inquire_inp.3g
/usr/man/man3.Z/dcpolygon.3g	/usr/man/man3.Z/inquire_req.3g
/usr/man/man3.Z/dcpolyline.3g	/usr/man/man3.Z/inquire_siz.3g
/usr/man/man3.Z/dcpolymarke.3g	/usr/man/man3.Z/inquire_tex.3g
/usr/man/man3.Z/dcrectangle.3g	/usr/man/man3.Z/intarc.3g
/usr/man/man3.Z/dctext.3g	/usr/man/man3.Z/intbitmap_p.3g
/usr/man/man3.Z/default_kno.3g	/usr/man/man3.Z/intbitmap_t.3g
/usr/man/man3.Z/define_colo.3g	/usr/man/man3.Z/intblock_mo.3g
/usr/man/man3.Z/define_rast.3g	/usr/man/man3.Z/intblock_re.3g
/usr/man/man3.Z/define_trim.3g	/usr/man/man3.Z/intblock_wr.3g
/usr/man/man3.Z/depth_cue.3g	/usr/man/man3.Z/intcharacte.3g
/usr/man/man3.Z/depth_cue_c.3g	/usr/man/man3.Z/intcircle.3g
/usr/man/man3.Z/depth_cue_r.3g	/usr/man/man3.Z/intclip_rec.3g
/usr/man/man3.Z/depth_indic.3g	/usr/man/man3.Z/intconcat_m.3g
/usr/man/man3.Z/designate_c.3g	/usr/man/man3.Z/intconcat_t.3g
/usr/man/man3.Z/disable_eve.3g	/usr/man/man3.Z/intdraw2d.3g
/usr/man/man3.Z/intecho_typ.3g	/usr/man/man3.Z/polymarker3.3g
/usr/man/man3.Z/intecho_upd.3g	/usr/man/man3.Z/polyrectang.3g
/usr/man/man3.Z/interior_st.3g	/usr/man/man3.Z/pop_matrix.3g
/usr/man/man3.Z/inthatch_sp.3g	/usr/man/man3.Z/pop_matrix2.3g
/usr/man/man3.Z/intinquire_.3g	/usr/man/man3.Z/pop_matrix3.3g
/usr/man/man3.Z/intline_end.3g	/usr/man/man3.Z/push_matrix.3g
/usr/man/man3.Z/intline_rep.3g	/usr/man/man3.Z/push_vdc.3g
/usr/man/man3.Z/intline_wid.3g	/usr/man/man3.Z/push_vdc_ma.3g
/usr/man/man3.Z/intmove2d.3g	/usr/man/man3.Z/quadrilater.3g
/usr/man/man3.Z/intpartial_.3g	/usr/man/man3.Z/read_choice.3g
/usr/man/man3.Z/intperimete.3g	/usr/man/man3.Z/read_locato.3g
/usr/man/man3.Z/intpolygon2.3g	/usr/man/man3.Z/rectangle.3g
/usr/man/man3.Z/intpolyline.3g	/usr/man/man3.Z/replace_mat.3g
/usr/man/man3.Z/intpop_matr.3g	/usr/man/man3.Z/request_choi.3g
/usr/man/man3.Z/intpush_mat.3g	/usr/man/man3.Z/request_loc.3g
/usr/man/man3.Z/intra_chara.3g	/usr/man/man3.Z/rgb_to_inde.3g
/usr/man/man3.Z/intrectangl.3g	/usr/man/man3.Z/sample_choi.3g
/usr/man/man3.Z/intreplace_.3g	/usr/man/man3.Z/sample_loca.3g
/usr/man/man3.Z/intrequest_.3g	/usr/man/man3.Z/set_hit_mod.3g

```

/usr/man/man3.Z/intsample_l.3g
/usr/man/man3.Z/intset_pick.3g
/usr/man/man3.Z/inttext2d.3g
/usr/man/man3.Z/inttext_ori.3g
/usr/man/man3.Z/inttransfor.3g
/usr/man/man3.Z/intvdc_exte.3g
/usr/man/man3.Z/intview_mat.3g
/usr/man/man3.Z/intview_por.3g
/usr/man/man3.Z/intview_win.3g
/usr/man/man3.Z/knot_vector.3g
/usr/man/man3.Z/light_ambie.3g
/usr/man/man3.Z/light_atten.3g
/usr/man/man3.Z/light_model.3g
/usr/man/man3.Z/light_sourc.3g
/usr/man/man3.Z/light_switc.3g
/usr/man/man3.Z/line_color.3g
/usr/man/man3.Z/line_color_.3g
/usr/man/man3.Z/line_endpoi.3g
/usr/man/man3.Z/line_repeat.3g
/usr/man/man3.Z/line_type.3g
/usr/man/man3.Z/line_width.3g
/usr/man/man3.Z/make_pictur.3g
/usr/man/man3.Z/mapping_mod.3g
/usr/man/man3.Z/marker_colo.3g
/usr/man/man3.Z/marker_orie.3g
/usr/man/man3.Z/marker_size.3g
/usr/man/man3.Z/marker_type.3g
/usr/man/man3.Z/move.3g
/usr/man/man3.Z/move2d.3g
/usr/man/man3.Z/move3d.3g
/usr/man/man3.Z/partial_arc.3g
/usr/man/man3.Z/partial_ell.3g
/usr/man/man3.Z/partial_pol.3g
/usr/man/man3.Z/pattern_def.3g
/usr/man/man3.Z/perimeter_c.3g
/usr/man/man3.Z/perimeter_r.3g
/usr/man/man3.Z/perimeter_t.3g
/usr/man/man3.Z/pick_depth.3g
/usr/man/man3.Z/pick_window.3g
/usr/man/man3.Z/polygon.3g
/usr/man/man3.Z/polygon2d.3g
/usr/man/man3.Z/polygon3d.3g
/usr/man/man3.Z/polyline.3g
/usr/man/man3.Z/polyline2d.3g
/usr/man/man3.Z/polyline3d.3g
/usr/man/man3.Z/set_locator.3g
/usr/man/man3.Z/set_pi_p2.3g
/usr/man/man3.Z/set_pick_de.3g
/usr/man/man3.Z/set_pick_wi.3g
/usr/man/man3.Z/set_signals.3g
/usr/man/man3.Z/shade_mode.3g
/usr/man/man3.Z/shade_range.3g
/usr/man/man3.Z/spline.3g
/usr/man/man3.Z/spline_curv.3g
/usr/man/man3.Z/spline_surf.3g
/usr/man/man3.Z/starbase.3g
/usr/man/man3.Z/surface_coe.3g
/usr/man/man3.Z/surface_mod.3g
/usr/man/man3.Z/text.3g
/usr/man/man3.Z/text2d.3g
/usr/man/man3.Z/text3d.3g
/usr/man/man3.Z/text_alignm.3g
/usr/man/man3.Z/text_color.3g
/usr/man/man3.Z/text_color_.3g
/usr/man/man3.Z/text_font_i.3g
/usr/man/man3.Z/text_line_p.3g
/usr/man/man3.Z/text_line_s.3g
/usr/man/man3.Z/text_orient.3g
/usr/man/man3.Z/text_path.3g
/usr/man/man3.Z/text_precis.3g
/usr/man/man3.Z/text_switch.3g
/usr/man/man3.Z/track.3g
/usr/man/man3.Z/track_off.3g
/usr/man/man3.Z/transform_p.3g
/usr/man/man3.Z/triangular_.3g
/usr/man/man3.Z/trimming_cu.3g
/usr/man/man3.Z/u_knot_vect.3g
/usr/man/man3.Z/v_knot_vect.3g
/usr/man/man3.Z/vdc_extent.3g
/usr/man/man3.Z/vdc_justifi.3g
/usr/man/man3.Z/vdc_to_dc.3g
/usr/man/man3.Z/vdc_to_wc.3g
/usr/man/man3.Z/vertex_form.3g
/usr/man/man3.Z/view_camera.3g
/usr/man/man3.Z/view_matrix.3g
/usr/man/man3.Z/view_port.3g
/usr/man/man3.Z/view_volume.3g
/usr/man/man3.Z/view_window.3g
/usr/man/man3.Z/viewpoint.3g
/usr/man/man3.Z/viewport_ju.3g

```

/usr/man/man3.Z/polymarker.3g	/usr/man/man3.Z/wc_to_vdc.3g
/usr/man/man3.Z/polymarker2.3g	/usr/man/man3.Z/write_enabl.3g
/usr/man/man3.Z/zbuffer_swi.3g	/usr/man/man4.Z/bitmapfile.4

## Fileset PFA\_FM

/system/PFA_FM/revlist	/usr/man/man3.Z/fainit.3w
/usr/include/fa.h	/usr/man/man3.Z/farectwrite.3w
/usr/include/fonticon.h	/usr/man/man3.Z/faroll.3w
/usr/lib/fa_fm_demo/Makefile	/usr/man/man3.Z/fasetinfo.3w
/usr/lib/fa_fm_demo/README	/usr/man/man3.Z/faterminate.3w
/usr/lib/fa_fm_demo/fontbyte.c	/usr/man/man3.Z/fawrite.3w
/usr/lib/fa_fm_demo/fontcurs.c	/usr/man/man3.Z/fm_activate.3w
/usr/lib/fa_fm_demo/fontdump.c	/usr/man/man3.Z/fm_clipflag.3w
/usr/lib/fa_fm_demo/fontinfo.c	/usr/man/man3.Z/fm_cliplim.3w
/usr/lib/fa_fm_demo/fontopt.c	/usr/man/man3.Z/fm_colors.3w
/usr/lib/fa_fm_demo/fontprop.c	/usr/man/man3.Z/fm_fileinfo.3w
/usr/lib/libfa.a	/usr/man/man3.Z/fm_fontdir.3w
/usr/lib/libfontm.a	/usr/man/man3.Z/fm_getname.3w
/usr/man/man3.Z/fa.3w	/usr/man/man3.Z/fm_load.3w
/usr/man/man3.Z/faclear.3w	/usr/man/man3.Z/fm_opt.3w
/usr/man/man3.Z/facolors.3w	/usr/man/man3.Z/fm_rasterin.3w
/usr/man/man3.Z/facursor.3w	/usr/man/man3.Z/fm_remove.3w
/usr/man/man3.Z/fafontactiv.3w	/usr/man/man3.Z/fm_str_len.3w
/usr/man/man3.Z/fafontload.3w	/usr/man/man3.Z/fm_styleinf.3w
/usr/man/man3.Z/fafontremov.3w	/usr/man/man3.Z/fm_write.3w
/usr/man/man3.Z/fagetinfo.3w	/usr/man/man3.Z/fontm.3w

## Fileset PSTARBAS

/system/PSTARBAS/CDFcustomize	/usr/lib/libgkssb.20.a
/system/PSTARBAS/customize	/usr/lib/libgkssb.a
/system/PSTARBAS/revlist	/usr/lib/libsb1.10.a
/usr/include/sdi.c.h	/usr/lib/libsb1.20.a
/usr/include/starbase.c.h	/usr/lib/libsb1.a
/usr/include/starbase.f1.h	/usr/lib/libsb2.a
/usr/include/starbase.f2.h	/usr/lib/libsbcmg.a
/usr/include/starbase.p1.h	/usr/lib/llib-lsb
/usr/include/starbase.p2.h	/usr/lib/llib-lsb.ln
/usr/lib/libXwindow.a	/usr/lib/starbase/formatters/fmt_table.c
/usr/lib/libdd262x.a	/usr/lib/starbase/formatters/pcl/bmprint.c
/usr/lib/libdd300h.a	/usr/lib/starbase/formatters/pcl/cfg.ctmplt
/usr/lib/libdd300l.a	/usr/lib/starbase/formatters/pcl/cfg.template
/usr/lib/libdd9836a.a	/usr/lib/starbase/formatters/pcl/libfmpcl.a
/usr/lib/libdd9836c.a	/usr/lib/starbase/lib98720w_d.a

/usr/lib/libdd9837.a	/usr/lib/starbase/libdd262x_d.a
/usr/lib/libdd98550.a	/usr/lib/starbase/libdd300h_d.a
/usr/lib/libdd98556.a	/usr/lib/starbase/libdd300l_d.a
/usr/lib/libdd98700.a	/usr/lib/starbase/libdd98550_d.a
/usr/lib/libdd98710.a	/usr/lib/starbase/libdd98556_d.a
/usr/lib/libdd98720.a	/usr/lib/starbase/libdd98720_d.a
/usr/lib/libdd98720w.a	/usr/lib/starbase/libdd98721_d.a
/usr/lib/libdd98721.a	/usr/lib/starbase/libdd98730_d.a
/usr/lib/libdd98730.a	/usr/lib/starbase/libdd98731_d.a
/usr/lib/libdd98731.a	/usr/lib/starbase/libddXn_d.a
/usr/lib/libddCADplt.a	/usr/lib/starbase/libddbit_d.a
/usr/lib/libddSMDpix.a	/usr/lib/starbase/libddbyte_d.a
/usr/lib/libddSMDpln.a	/usr/lib/starbase/libddhil_d.a
/usr/lib/libddXn.a	/usr/lib/starbase/libddhpgl_d.a
/usr/lib/libddbit.a	/usr/lib/starbase/libddlkbd_d.a
/usr/lib/libddbyte.a	/usr/lib/starbase/libddsox11_d.a
/usr/lib/libddgks.kbd.a	/usr/lib/starbase/libddterm_d.a
/usr/lib/libddgksm.a	/usr/lib/starbase/libddx11p_d.a
/usr/lib/libddhil.a	/usr/lib/starbase/sb_daemon.o
/usr/lib/libddhpcgm.a	/usr/lib/starbase/sddg/README
/usr/lib/libddhpgl.a	/usr/lib/starbase/sddg/body/sbdvr.o
/usr/lib/libddhpterm.a	/usr/lib/starbase/sddg/body/sbdvr_d.o
/usr/lib/libddkbd.a	/usr/lib/starbase/sddg/head/sbdvr.h
/usr/lib/libddlkbd.a	/usr/lib/starbase/sddg/head/sbdvr_typ.h
/usr/lib/libddsox11.a	/usr/lib/starbase/sddg/head/sbdvr_util.h
/usr/lib/libgkssb.10.a	/usr/lib/starbase/sddg/mybody/sbdvr_tab.c

## Partition X11WINDOWS

### Fileset PX11MAN

/system/PX11MAN/customize	/usr/man/man3.Z/XDrawText.3x
/system/PX11MAN/revlist	/usr/man/man3.Z/XEmptyRegio.3x
/usr/man/man1.Z/X.1	/usr/man/man3.Z/XFillRectan.3x
/usr/man/man1.Z/Xserver.1	/usr/man/man3.Z/XFlush.3x
/usr/man/man1.Z/bitmap.1	/usr/man/man3.Z/XFree.3x
/usr/man/man1.Z/gwindstop.1	/usr/man/man3.Z/XGetDefault.3x
/usr/man/man1.Z/hpterm.1	/usr/man/man3.Z/XGetVisualI.3x
/usr/man/man1.Z/hpwm.1	/usr/man/man3.Z/XGetWindowA.3x
/usr/man/man1.Z/resize.1	/usr/man/man3.Z/XGetWindowP.3x
/usr/man/man1.Z/rgb.1	/usr/man/man3.Z/XGrabButton.3x
/usr/man/man1.Z/sb2xwd.1	/usr/man/man3.Z/XGrabKey.3x
/usr/man/man1.Z/uwm.1	/usr/man/man3.Z/XGrabKeyboa.3x
/usr/man/man1.Z/x11start.1	/usr/man/man3.Z/XGrabPointe.3x

```

/usr/man/man1.Z/xclock.1
/usr/man/man1.Z/xfc.1
/usr/man/man1.Z/xfd.1
/usr/man/man1.Z/xhost.1
/usr/man/man1.Z/xinit.1
/usr/man/man1.Z/xinitcolorm.1
/usr/man/man1.Z/xload.1
/usr/man/man1.Z/xmodmap.1
/usr/man/man1.Z/xpr.1
/usr/man/man1.Z/xprkbd.1
/usr/man/man1.Z/xrdb.1
/usr/man/man1.Z/xrefresh.1
/usr/man/man1.Z/xseethru.1
/usr/man/man1.Z/xset.1
/usr/man/man1.Z/xsetroot.1
/usr/man/man1.Z/xterm.1
/usr/man/man1.Z/xwcreate.1
/usr/man/man1.Z/xwd.1
/usr/man/man1.Z/xwd2sb.1
/usr/man/man1.Z/xwdestroy.1
/usr/man/man1.Z/xwininfo.1
/usr/man/man1.Z/xwud.1
/usr/man/man3.Z/CONSTRAINT.3x
/usr/man/man3.Z/CORE.3x
/usr/man/man3.Z/Constraint.3x
/usr/man/man3.Z/Core.3x
/usr/man/man3.Z/Intro.3x
/usr/man/man3.Z/XAddHost.3x
/usr/man/man3.Z/XAllocColor.3x
/usr/man/man3.Z/XAllowEvent.3x
/usr/man/man3.Z/XChangeKeyC.3x
/usr/man/man3.Z/XChangeKeyM.3x
/usr/man/man3.Z/XChangePoin.3x
/usr/man/man3.Z/XChangeSave.3x
/usr/man/man3.Z/XChangeWind.3x
/usr/man/man3.Z/XClearArea.3x
/usr/man/man3.Z/XConfigureW.3x
/usr/man/man3.Z/XCopyArea.3x
/usr/man/man3.Z/XCreateColo.3x
/usr/man/man3.Z/XCreateFont.3x
/usr/man/man3.Z/XCreateGC.3x
/usr/man/man3.Z/XCreateImag.3x
/usr/man/man3.Z/XCreatePixm.3x
/usr/man/man3.Z/XGrabServer.3x
/usr/man/man3.Z/XHPAcknowledge.3x
/usr/man/man3.Z/XHPDeviceAu.3x
/usr/man/man3.Z/XHPDisableR.3x
/usr/man/man3.Z/XHPEnableRe.3x
/usr/man/man3.Z/XHPFreeDevi.3x
/usr/man/man3.Z/XHPGet16bit.3x
/usr/man/man3.Z/XHPGetCurre.3x
/usr/man/man3.Z/XHPGetDevic.3x
/usr/man/man3.Z/XHPGetExtEv.3x
/usr/man/man3.Z/XHPGetServe.3x
/usr/man/man3.Z/XHPGrabDevi.3x
/usr/man/man3.Z/XHPis16bitC.3x
/usr/man/man3.Z/XHPKeysymTo.3x
/usr/man/man3.Z/XHPListInpu.3x
/usr/man/man3.Z/XHPPrompt.3x
/usr/man/man3.Z/XHPRoman8To.3x
/usr/man/man3.Z/XHPSelectEx.3x
/usr/man/man3.Z/XHPSetDevic.3x
/usr/man/man3.Z/XHPsetError.3x
/usr/man/man3.Z/XHPSetEvent.3x
/usr/man/man3.Z/XHPSetInput.3x
/usr/man/man3.Z/XHPUngrabDe.3x
/usr/man/man3.Z/XIfEvent.3x
/usr/man/man3.Z/XInstallCol.3x
/usr/man/man3.Z/XInternAtom.3x
/usr/man/man3.Z/XIntersectR.3x
/usr/man/man3.Z/XListFonts.3x
/usr/man/man3.Z/XLoadFont.3x
/usr/man/man3.Z/XLookupKeys.3x
/usr/man/man3.Z/XMapWindow.3x
/usr/man/man3.Z/XOpenDispla.3x
/usr/man/man3.Z/XParseGeome.3x
/usr/man/man3.Z/XPolygonReg.3x
/usr/man/man3.Z/XPutBackEve.3x
/usr/man/man3.Z/XPutImage.3x
/usr/man/man3.Z/XQueryBestS.3x
/usr/man/man3.Z/XQueryColor.3x
/usr/man/man3.Z/XQueryPoint.3x
/usr/man/man3.Z/XQueryTree.3x
/usr/man/man3.Z/XRaiseWindo.3x
/usr/man/man3.Z/XReadBitmap.3x
/usr/man/man3.Z/XRecolorCur.3x

```

/usr/man/man3.Z/XCreateRegi.3x	/usr/man/man3.Z/XReparentWi.3x
/usr/man/man3.Z/XCreateWind.3x	/usr/man/man3.Z/XSaveContex.3x
/usr/man/man3.Z/XDefineCurs.3x	/usr/man/man3.Z/XSelectInpu.3x
/usr/man/man3.Z/XDestroyWin.3x	/usr/man/man3.Z/XSetArcMode.3x
/usr/man/man3.Z/XDrawArc.3x	/usr/man/man3.Z/XSetClassHi.3x
/usr/man/man3.Z/XDrawImageS.3x	/usr/man/man3.Z/XSetClipOri.3x
/usr/man/man3.Z/XDrawLine.3x	/usr/man/man3.Z/XSetCloseDo.3x
/usr/man/man3.Z/XDrawPoint.3x	/usr/man/man3.Z/XSetCommand.3x
/usr/man/man3.Z/XDrawRectan.3x	/usr/man/man3.Z/XsetErrorHa.3x
/usr/man/man3.Z/XDrawString.3x	/usr/man/man3.Z/XSetFillSty.3x
/usr/man/man3.Z/XSetFont.3x	/usr/man/man3.Z/XrEditorGro.3x
/usr/man/man3.Z/XSetFontPat.3x	/usr/man/man3.Z/XrErrno.3x
/usr/man/man3.Z/XSetIconNam.3x	/usr/man/man3.Z/XrGlobals.3x
/usr/man/man3.Z/XSetIconSiz.3x	/usr/man/man3.Z/XrGroupBox.3x
/usr/man/man3.Z/XSetInputFo.3x	/usr/man/man3.Z/XrInit.3x
/usr/man/man3.Z/XSetLineAtt.3x	/usr/man/man3.Z/XrInput.3x
/usr/man/man3.Z/XSetNormalH.3x	/usr/man/man3.Z/XrInputCo16.3x
/usr/man/man3.Z/XSetPointer.3x	/usr/man/man3.Z/XrInputConf.3x
/usr/man/man3.Z/XSetScreenS.3x	/usr/man/man3.Z/XrInputInit.3x
/usr/man/man3.Z/XSetSelecti.3x	/usr/man/man3.Z/XrInputMap.3x
/usr/man/man3.Z/XSetSizeHin.3x	/usr/man/man3.Z/XrInputMute.3x
/usr/man/man3.Z/XSetStandC.3x	/usr/man/man3.Z/XrListEdit.3x
/usr/man/man3.Z/XSetStandP.3x	/usr/man/man3.Z/XrMenu.3x
/usr/man/man3.Z/XSetState.3x	/usr/man/man3.Z/XrMessageB.3x
/usr/man/man3.Z/XSetTile.3x	/usr/man/man3.Z/XrPageEdit.3x
/usr/man/man3.Z/XSetTransie.3x	/usr/man/man3.Z/XrPanel.3x
/usr/man/man3.Z/XSetWMHints.3x	/usr/man/man3.Z/XrPoint.3x
/usr/man/man3.Z/XSetZoomHin.3x	/usr/man/man3.Z/XrPushButto.3x
/usr/man/man3.Z/XStoreBytes.3x	/usr/man/man3.Z/XrRadioButt.3x
/usr/man/man3.Z/XStoreColor.3x	/usr/man/man3.Z/XrRasterEdi.3x
/usr/man/man3.Z/XStoreName.3x	/usr/man/man3.Z/XrRasterSel.3x
/usr/man/man3.Z/XStringToKe.3x	/usr/man/man3.Z/XrRectangle.3x
/usr/man/man3.Z/XSynchroniz.3x	/usr/man/man3.Z/XrResource.3x
/usr/man/man3.Z/XTextExtent.3x	/usr/man/man3.Z/XrScrollBar.3x
/usr/man/man3.Z/XTextWidth.3x	/usr/man/man3.Z/XrStaticRas.3x
/usr/man/man3.Z/XTranslateC.3x	/usr/man/man3.Z/XrStaticTex.3x
/usr/man/man3.Z/XUnmapWindo.3x	/usr/man/man3.Z/XrTextEdit.3x
/usr/man/man3.Z/XWARROW.3x	/usr/man/man3.Z/XrTitleBar.3x

/usr/man/man3.Z/XWBULLETIN.3x	/usr/man/man3.Z/XrUtilities.3x
/usr/man/man3.Z/XWBUTTON.3x	/usr/man/man3.Z/XwArrow.3x
/usr/man/man3.Z/XWCASCADE.3x	/usr/man/man3.Z/XwBBoard.3x
/usr/man/man3.Z/XWCREATETIL.3x	/usr/man/man3.Z/XwButton.3x
/usr/man/man3.Z/XWFORM.3x	/usr/man/man3.Z/XwCascade.3x
/usr/man/man3.Z/XWFRAME.3x	/usr/man/man3.Z/XwCreateTil.3x
/usr/man/man3.Z/XWIMAGEEDIT.3x	/usr/man/man3.Z/XwForm.3x
/usr/man/man3.Z/XWLIST.3x	/usr/man/man3.Z/XwFrame.3x
/usr/man/man3.Z/XWMANAGER.3x	/usr/man/man3.Z/XwImageEdit.3x
/usr/man/man3.Z/XWMENUBTN.3x	/usr/man/man3.Z/XwList.3x
/usr/man/man3.Z/XWMENUPANE.3x	/usr/man/man3.Z/XwManager.3x
/usr/man/man3.Z/XWMENUSEP.3x	/usr/man/man3.Z/XwMenuBtn.3x
/usr/man/man3.Z/XWMOVEFOCUS.3x	/usr/man/man3.Z/XwMenuMgr.3x
/usr/man/man3.Z/XWPANEL.3x	/usr/man/man3.Z/XwMenuPane.3x
/usr/man/man3.Z/XWPOPUPMGR.3x	/usr/man/man3.Z/XwMenuSep.3x
/usr/man/man3.Z/XWPRIMITIVE.3x	/usr/man/man3.Z/XwMoveFocus.3x
/usr/man/man3.Z/XWPULLDOWN.3x	/usr/man/man3.Z/XwPButton.3x
/usr/man/man3.Z/XWPUSHBUTTO.3x	/usr/man/man3.Z/XwPanel.3x
/usr/man/man3.Z/XWREGISTERC.3x	/usr/man/man3.Z/XwPopupMgr.3x
/usr/man/man3.Z/XWROWCOLUMN.3x	/usr/man/man3.Z/XwPrimitive.3x
/usr/man/man3.Z/XWSASH.3x	/usr/man/man3.Z/XwPulldown.3x
/usr/man/man3.Z/XWSCROLLBAR.3x	/usr/man/man3.Z/XwRManager.3x
/usr/man/man3.Z/XWSCROLLEDW.3x	/usr/man/man3.Z/XwRegisterC.3x
/usr/man/man3.Z/XWSTATICRAS.3x	/usr/man/man3.Z/XwSRaster.3x
/usr/man/man3.Z/XWSTATICTEX.3x	/usr/man/man3.Z/XwSText.3x
/usr/man/man3.Z/XWTEXTEDIT.3x	/usr/man/man3.Z/XwSWindow.3x
/usr/man/man3.Z/XWTITLEBAR.3x	/usr/man/man3.Z/XwSash.3x
/usr/man/man3.Z/XWTOGGLE.3x	/usr/man/man3.Z/XwScrollBar.3x
/usr/man/man3.Z/XWVALUATOR.3x	/usr/man/man3.Z/XwTextEdit.3x
/usr/man/man3.Z/XWVPANED.3x	/usr/man/man3.Z/XwTitleBar.3x
/usr/man/man3.Z/XWWORKSPACE.3x	/usr/man/man3.Z/XwToggle.3x
/usr/man/man3.Z/XWarpPointe.3x	/usr/man/man3.Z/XwVPW.3x
/usr/man/man3.Z/XWindowEven.3x	/usr/man/man3.Z/XwValuator.3x
/usr/man/man3.Z/XrCheckBox.3x	/usr/man/man3.Z/XwWorkSpace.3x
/usr/man/man3.Z/XrEditor.3x	/usr/man/man4.Z/bdf.4

## Fileset PX11PENV

/system/PX11PENV/customize	/usr/include/X11/Shell.h
/system/PX11PENV/revlist	/usr/include/X11/ShellP.h
/usr/include/X11/StringDefs.h	/usr/include/X11/TM.h
/usr/include/X11/Vendor.h	/usr/include/X11/VendorP.h
/usr/include/X11/X.h	/usr/include/X11/X10.h
/usr/include/X11/XHPImageIO.h	/usr/include/X11/XHPLib.h
/usr/include/X11/XHPproto.h	/usr/include/X11/XWDFFile.h



```

/usr/include/X11/Xatom.h
/usr/include/X11/Xmd.h
/usr/include/X11/Xproto.h
/usr/include/X11/Xresource.h
/usr/contrib/Xf11/Makefile
/usr/contrib/Xf11/sample2.f
/usr/include/X11/bezier.h
/usr/include/X11/cursorfont.h
/usr/include/X11/keysymdef.h
/usr/include/Xf11/Xflib.h
/usr/include/Xf11/xfdeclare.h
/usr/include/Xf11/xffields.h
/usr/include/Xf11/xfatypes.h
/usr/include/Xr11/Xrlib.h
/usr/include/Xr11/in_types.h
/usr/include/Xr11/types.h
/usr/include/Xw/ArrowP.h
/usr/include/Xw/BBoardP.h
/usr/include/Xw/CascadeP.h
/usr/include/Xw/Form.h
/usr/include/Xw/ImageEdit.h
/usr/contrib/Xw/examples/README
/usr/contrib/Xw/examples/demo1.c
/usr/contrib/Xw/examples/demo2.c
/usr/include/X11/Cardinals.h
/usr/include/X11/Composite.h
/usr/include/X11/CompositeP.h
/usr/include/X11/ConstrainP.h
/usr/include/X11/Constraint.h
/usr/include/X11/Convert.h
/usr/include/X11/Core.h
/usr/include/X11/CoreP.h
/usr/include/X11/DECkeysym.h
/usr/include/X11/Intrinsic.h
/usr/include/X11/IntrinsicP.h
/usr/include/X11/Misc.h
/usr/include/X11/Xlib.h
/usr/include/X11/Xos.h
/usr/include/X11/Xprotostr.h
/usr/include/X11/Xutil.h
/usr/contrib/Xf11/sample1.f
/usr/contrib/Xf11/sample3.f
/usr/include/X11/copyright.h
/usr/include/X11/keysym.h
/usr/include/X11/zoid.h
/usr/include/Xf11/xfalias.h
/usr/include/Xf11/xfdefines.h
/usr/include/Xf11/xfmsg.h
/usr/include/Xr11/XGetKbdId.h
/usr/include/Xr11/defs.h
/usr/include/Xr11/keycode.h
/usr/include/Xw/Arrow.h
/usr/include/Xw/BBoard.h
/usr/include/Xw/Cascade.h
/usr/include/Xw/DisplayP.h
/usr/include/Xw/List.h
/usr/include/Xw/ImageEditP.h
/usr/include/Xw/FormP.h
/usr/include/Xw/Frame.h
/usr/include/Xw/FrameP.h
/usr/include/Xw/ListP.h
/usr/include/Xw/ListRC.h
/usr/include/Xw/ListRCP.h
/usr/include/Xw/MenuBtn.h
/usr/include/Xw/MenuBtnP.h
/usr/include/Xw/MenuPDBtn.h
/usr/include/Xw/MenuPDBtnP.h
/usr/include/Xw/MenuSep.h
/usr/include/Xw/MenuSepP.h
/usr/include/Xw/PButton.h
/usr/include/Xw/PButtonP.h
/usr/include/Xw/Panel.h

```

```

/usr/include/X11/Quarks.h           /usr/include/Xw/PanelP.h
/usr/include/Xw/PopupMgr.h         /usr/include/Xw/VPWP.h
/usr/include/Xw/PopupMgrP.h       /usr/include/Xw/Valuator.h
/usr/include/Xw/Pulldown.h         /usr/include/Xw/ValuatorP.h
/usr/include/Xw/PulldownP.h       /usr/include/Xw/WClassing.h
/usr/include/Xw/RcManager.h        /usr/include/Xw/WClassingP.h
/usr/include/Xw/RcManagerP.h      /usr/include/Xw/WorkSpace.h
/usr/include/Xw/SRaster.h          /usr/include/Xw/WorkSpaceP.h
/usr/include/Xw/SRasterP.h        /usr/include/Xw/Xw.h
/usr/include/Xw/SText.h           /usr/include/Xw/XwP.h
/usr/include/Xw/STextP.h          /usr/lib/libX11.a
/usr/include/Xw/SWindow.h         /usr/lib/libXext.a
/usr/include/Xw/SWindowP.h        /usr/lib/libXf11.a
/usr/include/Xw/Sash.h            /usr/lib/libXhp11.a
/usr/include/Xw/SashP.h           /usr/lib/libXr11.a
/usr/include/Xw/ScrollBar.h       /usr/lib/libXt.a
/usr/include/Xw/ScrollBarP.h      /usr/lib/libXw.a
/usr/include/Xw/SourceP.h         /usr/lib/liboldX.a
/usr/include/Xw/TextEdit.h        /usr/lib/llib-lX11.ln
/usr/include/Xw/TextEditP.h      /usr/lib/llib-lXext.ln
/usr/include/Xw/TitleBar.h        /usr/lib/llib-lXhp11.ln
/usr/include/Xw/TitleBarP.h      /usr/lib/llib-lXt.ln
/usr/include/Xw/Toggle.h          /usr/lib/llib-lXw.ln
/usr/include/Xw/ToggleP.h        /usr/lib/llib-loldX.ln
/usr/include/Xw/VPW.h
/usr/contrib/Xf11/HelloWorld/DrawSpec.h
/usr/contrib/Xf11/HelloWorld/DrawVal.h
/usr/contrib/Xf11/HelloWorld/FrDecl.h
/usr/contrib/Xf11/HelloWorld/FrSpec.h
/usr/contrib/Xf11/HelloWorld/FrVal.h
/usr/contrib/Xf11/HelloWorld/HelloSpec.h
/usr/contrib/Xf11/HelloWorld/HelloVal.h
/usr/contrib/Xf11/HelloWorld/HelloWorld.h
/usr/contrib/Xf11/HelloWorld/Makefile
/usr/contrib/Xf11/HelloWorld/SmDecl.h
/usr/contrib/Xf11/HelloWorld/SmSpec.h
/usr/contrib/Xf11/HelloWorld/SmVal.h
/usr/contrib/Xf11/HelloWorld/sample4.f
/usr/contrib/Xr11/examples/HelloWorld/HelloWorld.h
/usr/contrib/Xr11/examples/HelloWorld/Makefile
/usr/contrib/Xr11/examples/HelloWorld/frown.h
/usr/contrib/Xr11/examples/HelloWorld/sample1.c

```

```

/usr/contrib/Xr11/examples/HelloWorld/sample2.c
/usr/contrib/Xr11/examples/HelloWorld/sample3.c
/usr/contrib/Xr11/examples/HelloWorld/sample4.c
/usr/contrib/Xr11/examples/HelloWorld/smile.h
/usr/contrib/Xr11/examples/RB/Makefile
/usr/contrib/Xr11/examples/RB/buildPanel.c
/usr/contrib/Xr11/examples/RB/buildPanel.h
/usr/contrib/Xr11/examples/RB/file.c
/usr/contrib/Xr11/examples/RB/filePanel.c
/usr/contrib/Xr11/examples/RB/filePanel.h
/usr/contrib/Xr11/examples/RB/image.c
/usr/contrib/Xr11/examples/RB/image.h
/usr/contrib/Xr11/examples/RB/main.c
/usr/contrib/Xr11/examples/RB/main.h
/usr/contrib/Xr11/examples/RB/selRaster.c
/usr/contrib/Xw/examples/DEmo
/usr/contrib/Xw/examples/Makefile
/usr/contrib/Xw/examples/MultiButton/MButton.c
/usr/contrib/Xw/examples/MultiButton/MButton.h
/usr/contrib/Xw/examples/MultiButton/MButtonP.h
/usr/contrib/Xw/examples/MultiButton/README
/usr/contrib/Xw/examples/MultiButton/mButton
/usr/contrib/Xw/examples/MultiButton/mTest.make
/usr/contrib/Xw/examples/MultiButton/multiTest.c
/usr/contrib/Xw/examples/menudemo.c
/usr/contrib/Xw/examples/stfont.c
/usr/contrib/Xw/examples/tilettest.c

```

## Partition WINDOW

### Fileset DWINDOW

```

/system/DWINDOW/customize           /system/DWINDOW/revlist
/usr/lib/hpwindows/demosrc/Makefile  /usr/lib/hpwindows/demosrc/objonly/calc
/usr/lib/hpwindows/demosrc/README    /usr/lib/hpwindows/demosrc/repaint.c
/usr/lib/hpwindows/demosrc/Wfont     /usr/lib/hpwindows/demosrc/seticon.c
/usr/lib/hpwindows/demosrc/chcolor   /usr/lib/hpwindows/demosrc/sprite.c
/usr/lib/hpwindows/demosrc/crabs.c   /usr/lib/hpwindows/demosrc/wcal
/usr/lib/hpwindows/demosrc/domenu.c  /usr/lib/hpwindows/demosrc/wclock.c
/usr/lib/hpwindows/demosrc/fclock.f  /usr/lib/hpwindows/demosrc/wqix.c
/usr/lib/hpwindows/demosrc/lens.c    /usr/lib/hpwindows/demo/.HPCalc.msg
/usr/lib/hpwindows/demosrc/life.c    /usr/lib/hpwindows/demo/calc
/usr/lib/raster/12x16/calcfont       /usr/lib/hpwindows/demo/.HPCalc.hlp
/usr/lib/raster/icons/calculator     /usr/lib/raster/icons/mailbox_full

```

```

/usr/lib/hpwindows/demosrc/zot.1      /usr/lib/hpwindows/demosrc/wmdiag.c
/usr/lib/raster/6x8/calcfont          /usr/lib/hpwindows/demosrc/zot.c
/usr/lib/hpwindows/demosrc/editicon.c
/usr/lib/hpwindows/demosrc/fontbyte.c
/usr/lib/hpwindows/demosrc/fontcurs.c
/usr/lib/hpwindows/demosrc/fontdump.c
/usr/lib/hpwindows/demosrc/fontedit.c
/usr/lib/hpwindows/demosrc/fontinfo.c
/usr/lib/hpwindows/demosrc/fontopt.c
/usr/lib/raster/icons/mailbox_empty
/usr/lib/hpwindows/demosrc/fontprop.c
/usr/lib/hpwindows/demosrc/fontview.c
/usr/lib/hpwindows/demosrc/objonly/.HPCalc.hlp
/usr/lib/hpwindows/demosrc/objonly/.HPCalc.msg

```

## Fileset MWINDOW

```

/system/MWINDOW/customize            /usr/man/man3.Z/wautotop.3w
/system/MWINDOW/revlist              /usr/man/man3.Z/wbanner.3w
/usr/man/man3.Z/wbottom.3w           /usr/man/man3.Z/wconceal.3w
/usr/man/man3.Z/wcreate_gra.3w       /usr/man/man3.Z/wcreate_ter.3w
/usr/man/man3.Z/wdestroy.3w          /usr/man/man3.Z/wdfiltpos.3w
/usr/man/man3.Z/weventclear.3w       /usr/man/man3.Z/weventpoll.3w
/usr/man/man3.Z/wget_hw_spr.3w       /usr/man/man3.Z/wget_see_th.3w
/usr/man/man3.Z/wgetbcolor.3w       /usr/man/man3.Z/wgetbcoords.3w
/usr/man/man3.Z/wgetcoords.3w        /usr/man/man3.Z/wgetecho.3w
/usr/man/man3.Z/wgeticonpos.3w       /usr/man/man3.Z/wgetlocator.3w
/usr/man/man3.Z/wgetname.3w          /usr/man/man3.Z/wgetrastere.3w
/usr/man/man3.Z/wgetscreen.3w        /usr/man/man3.Z/wgetsigmask.3w
/usr/man/man3.Z/wgskbd.3w            /usr/man/man3.Z/whotspot_cr.3w
/usr/man/man3.Z/whotspot_de.3w       /usr/man/man3.Z/whotspot_ge.3w
/usr/man/man3.Z/whotspot_se.3w       /usr/man/man3.Z/wiconic.3w
/usr/man/man3.Z/window.3w            /usr/man/man3.Z/winit.3w
/usr/man/man3.Z/winput_conf.3w        /usr/man/man3.Z/winput_getr.3w
/usr/man/man3.Z/winput_read.3w       /usr/man/man3.Z/winput_setr.3w
/usr/man/man3.Z/winput_widp.3w       /usr/man/man3.Z/wmenu_activ.3w
/usr/man/man3.Z/wmenu_creat.3w       /usr/man/man3.Z/wmenu_delet.3w
/usr/man/man3.Z/wmenu_event.3w       /usr/man/man3.Z/wmenu_item.3w
/usr/man/man1.Z/wborder.1             /usr/man/man3.Z/wminquire.3w
/usr/man/man1.Z/wcreate.1             /usr/man/man3.Z/wmkill.3w
/usr/man/man1.Z/wdestroy.1           /usr/man/man3.Z/wmove.3w
/usr/man/man1.Z/wdisp.1              /usr/man/man3.Z/wmpathmake.3w
/usr/man/man1.Z/wfont.1              /usr/man/man3.Z/wmrepaint.3w
/usr/man/man1.Z/windows.1            /usr/man/man3.Z/wpan.3w
/usr/man/man1.Z/wlist.1              /usr/man/man3.Z/wpauseoutpu.3w

```

```

/usr/man/man1.Z/wmmove.1          /usr/man/man3.Z/wrecover.3w
/usr/man/man1.Z/wmready.1        /usr/man/man3.Z/wscrn_sprit.3w
/usr/man/man1.Z/wmstart.1        /usr/man/man3.Z/wscroll_get.3w
/usr/man/man1.Z/wmstop.1         /usr/man/man3.Z/wscroll_set.3w
/usr/man/man1.Z/wselect.1        /usr/man/man3.Z/wselect.3w
/usr/man/man1.Z/wsh.1            /usr/man/man3.Z/wset_hw_spr.3w
/usr/man/man1.Z/wsize.1          /usr/man/man3.Z/wset_see_th.3w
/usr/man/man3.Z/altfont_ter.3w   /usr/man/man3.Z/wsetbcolor.3w
/usr/man/man3.Z/basefont_te.3w   /usr/man/man3.Z/wsetecho.3w
/usr/man/man3.Z/fontgetid_t.3w   /usr/man/man3.Z/wseticon.3w
/usr/man/man3.Z/fontgetname.3w   /usr/man/man3.Z/wseticonpos.3w
/usr/man/man3.Z/fontload_te.3w   /usr/man/man3.Z/wsetlabel.3w
/usr/man/man3.Z/fontreplace.3w   /usr/man/man3.Z/wsetlocator.3w
/usr/man/man3.Z/fontsize_te.3w   /usr/man/man3.Z/wsetrastere.3w
/usr/man/man3.Z/fontswap_te.3w   /usr/man/man3.Z/wsetsigmask.3w
/usr/man/man3.Z/fromxy_term.3w   /usr/man/man3.Z/wsfk_mode.3w
/usr/man/man3.Z/toxy_term0.3w    /usr/man/man3.Z/wsfk_prog.3w
/usr/man/man3.Z/wautodestro.3w   /usr/man/man3.Z/wshuffle.3w
/usr/man/man3.Z/wautoselect.3w   /usr/man/man3.Z/wsize.3w
/usr/man/man3.Z/wterminate.3w    /usr/man/man3.Z/wtop.3w
/usr/lib/hpwindows/man_examples/basalt.c
/usr/lib/hpwindows/man_examples/base_load.c
/usr/lib/hpwindows/man_examples/btest.c
/usr/lib/hpwindows/man_examples/build_icon.c
/usr/lib/hpwindows/man_examples/clear_gr.c
/usr/lib/hpwindows/man_examples/conceal_t0.c
/usr/lib/hpwindows/man_examples/create_gr.c
/usr/lib/hpwindows/man_examples/create_t0.c
/usr/lib/hpwindows/man_examples/echo_hand.c
/usr/lib/hpwindows/man_examples/est_gr.c
/usr/lib/hpwindows/man_examples/est_t0.c
/usr/lib/hpwindows/man_examples/est_wm_com.c
/usr/lib/hpwindows/man_examples/getfontinfo.c
/usr/lib/hpwindows/man_examples/invert_bc.c
/usr/lib/hpwindows/man_examples/kill_wm.c
/usr/lib/hpwindows/man_examples/loc_in_user.c
/usr/lib/hpwindows/man_examples/pan_gr.c
/usr/lib/hpwindows/man_examples/pause_resume.c
/usr/lib/hpwindows/man_examples/poll_events.c
/usr/lib/hpwindows/man_examples/quarter_clip.c
/usr/lib/hpwindows/man_examples/replace_icon.c
/usr/lib/hpwindows/man_examples/reset_loc.c
/usr/lib/hpwindows/man_examples/rm_window.c

```

```

/usr/lib/hpwindows/man_examples/set_gr_labs.c
/usr/lib/hpwindows/man_examples/setlabel_gr.c
/usr/lib/hpwindows/man_examples/shrink_it.c
/usr/lib/hpwindows/man_examples/shrink_t0.c
/usr/lib/hpwindows/man_examples/shuffle_dn.c
/usr/lib/hpwindows/man_examples/stair_step.c
/usr/lib/hpwindows/man_examples/stretch_gr.c
/usr/lib/hpwindows/man_examples/term_gr.c
/usr/lib/hpwindows/man_examples/term_t0.c
/usr/lib/hpwindows/man_examples/term_wm_com.c
/usr/lib/hpwindows/man_examples/toggle_icon.c
/usr/lib/hpwindows/man_examples/toggle_sel.c
/usr/lib/hpwindows/man_examples/wbanner_sub.c
/usr/lib/hpwindows/man_examples/wrepaint.c
/usr/lib/hpwindows/man_examples/write_dn.c

```

## Fileset PWINDOW

```

/system/PWINDOW/customize      /usr/include/window.h
/system/PWINDOW/revlist       /usr/lib/libwindow.a

```

## Partition PROG\_LANGS

### Fileset PC

```

/bin/gprof      /usr/include/msgcat.h
/bin/prof       /usr/include/ndbm.h
/bin/time       /usr/include/ndir.h
/lib/c1         /usr/include/nl_ctype.h
/lib/c220      /usr/include/nl_types.h
/lib/ccom20    /usr/include/nlist.h
/lib/cpass1    /usr/include/oldfatal.h
/lib/cpass2    /usr/include/oldmacros.h
/lib/cpp.ansi  /usr/include/port.h
/system/PC/revlist /usr/include/pwd.h
/usr/bin/cb     /usr/include/ranlib.h
/usr/bin/cflow /usr/include/regexp.h
/usr/bin/cxref  /usr/include/rje.h
/usr/bin/lint   /usr/include/search.h
/usr/include/a.exec.h /usr/include/setjmp.h
/usr/include/a.out.h /usr/include/sgtty.h
/usr/include/acllib.h /usr/include/signal.h
/usr/include/alarm.h /usr/include/stand.h
/usr/include/ar.h   /usr/include/stdio.h

```

/usr/include/assert.h  
/usr/include/bsdtty.h  
/usr/include/checklist.h  
/usr/include/cluster.h  
/usr/include/core.h  
/usr/include/ctype.h  
/usr/include/curses.h  
/usr/include/cvtnum.h  
/usr/include/dbm.h  
/usr/include/dial.h  
/usr/include/dirent.h  
/usr/include/disktab.h  
/usr/include/dumprestor.h

/usr/include/string.h  
/usr/include/symbol.h  
/usr/include/sys/ac.h  
/usr/include/sys/acct.h  
/usr/include/sys/acl.h  
/usr/include/sys/amigo.h  
/usr/include/sys/audit.h  
/usr/include/sys/bootrom.h  
/usr/include/sys/bsdtty.h  
/usr/include/sys/buf.h  
/usr/include/sys/callout.h  
/usr/include/sys/cct.h  
/usr/include/sys/ciper.h

/usr/include/dvio.h	/usr/include/sys/clock.h
/usr/include/errnet.h	/usr/include/sys/cmap.h
/usr/include/errno.h	/usr/include/sys/conf.h
/usr/include/execargs.h	/usr/include/sys/cons.h
/usr/include/fatal.h	/usr/include/sys/cs80.h
/usr/include/fbackup.h	/usr/include/sys/dbg.h
/usr/include/fcntl.h	/usr/include/sys/dil.h
/usr/include/ftw.h	/usr/include/sys/dilio.h
/usr/include/globaldefs.h	/usr/include/sys/dir.h
/usr/include/grp.h	/usr/include/sys/dirent.h
/usr/include/iovalid.h	/usr/include/sys/dk.h
/usr/include/kern_prof.h	/usr/include/sys/dma.h
/usr/include/langinfo.h	/usr/include/sys/dmap.h
/usr/include/limits.h	/usr/include/sys/dmmsgtype.h
/usr/include/machine/clock.h	/usr/include/sys/dux_mbuf.h
/usr/include/machine/param.h	/usr/include/sys/errno.h
/usr/include/machine/pcb.h	/usr/include/sys/file.h
/usr/include/machine/psl.h	/usr/include/sys/filsys.h
/usr/include/machine/pte.h	/usr/include/sys/fs.h
/usr/include/machine/reg.h	/usr/include/sys/getaccess.h
/usr/include/machine/sendsig.h	/usr/include/sys/gpio.h
/usr/include/machine/trap.h	/usr/include/sys/graf.h
/usr/include/machine/vmparam.h	/usr/include/sys/graphics.h
/usr/include/macros.h	/usr/include/sys/hil.h
/usr/include/magic.h	/usr/include/sys/hilioctl.h
/usr/include/malloc.h	/usr/include/sys/hpib.h
/usr/include/math.h	/usr/include/sys/hpibio.h
/usr/include/memory.h	/usr/include/sys/ino.h
/usr/include/misc.h	/usr/include/sys/inode.h
/usr/include/mntent.h	/usr/include/sys/intrpt.h
/usr/include/mnttab.h	/usr/include/sys/iobuf.h
/usr/include/model.h	/usr/include/sys/ioctl.h
/usr/include/mon.h	/usr/include/sys/iomap.h
/usr/include/msgbuf.h	/usr/include/sys/ipc.h
/usr/include/sys/ite.h	/usr/include/sys/ti9914.h
/usr/include/sys/kbd.h	/usr/include/sys/time.h
/usr/include/sys/kbd_chars.h	/usr/include/sys/timeb.h
/usr/include/sys/kernel.h	/usr/include/sys/timeout.h
/usr/include/sys/lanfunc.h	/usr/include/sys/times.h
/usr/include/sys/lock.h	/usr/include/sys/trace.h
/usr/include/sys/lp.h	/usr/include/sys/trap.h
/usr/include/sys/lprio.h	/usr/include/sys/tryrec.h
/usr/include/sys/map.h	/usr/include/sys/ttold.h



```

/usr/include/sys/mf.h
/usr/include/sys/mknod.h
/usr/include/sys/mman.h
/usr/include/sys/modem.h
/usr/include/sys/mount.h
/usr/include/sys/msg.h
/usr/include/sys/msgbuf.h
/usr/include/sys/mtio.h
/usr/include/sys/mux.h
/usr/include/sys/nami.h
/usr/include/sys/nsp.h
/usr/include/sys/opt.h
/usr/include/sys/optflag.h
/usr/include/sys/param.h
/usr/include/sys/pcb.h
/usr/include/sys/pdi.h
/usr/include/sys/privgrp.h
/usr/include/sys/proc.h
/usr/include/sys/protocol.h
/usr/include/sys/psl.h
/usr/include/sys/pt.h
/usr/include/sys/ptrace.h
/usr/include/sys/pty.h
/usr/include/sys/ptyio.h
/usr/include/sys/reboot.h
/usr/include/sys/reg.h
/usr/include/sys/resource.h
/usr/include/sys/rfa_user.h
/usr/include/sys/rje_ioctl.h
/usr/include/sys/rjedriver.h
/usr/include/sys/rmswap.h
/usr/include/sys/rtprio.h
/usr/include/sys/scsi.h
/usr/include/sys/seg.h
/usr/include/sys/sem.h
/usr/include/sys/shm.h
/usr/include/sys/signal.h
/usr/include/sys/simon.h
/usr/include/sys/sitemap.h
/usr/include/sys/space.h
/usr/include/sys/srm.h
/usr/include/sys/stat.h
/usr/include/sys/sysmacros.h
/usr/include/sys/systm.h
/usr/include/sys/termio.h

/usr/include/sys/tty.h
/usr/include/sys/types.h
/usr/include/sys/uo.h
/usr/include/sys/unsp.h
/usr/include/sys/user.h
/usr/include/sys/utsname.h
/usr/include/sys/vfs.h
/usr/include/sys/vm.h
/usr/include/sys/vme.h
/usr/include/sys/vme2.h
/usr/include/sys/vmmac.h
/usr/include/sys/vmmeter.h
/usr/include/sys/vmparam.h
/usr/include/sys/vmsystem.h
/usr/include/sys/vnode.h
/usr/include/sys/vtest.h
/usr/include/sys/vtimes.h
/usr/include/sys/wait.h
/usr/include/tar.h
/usr/include/tcio.h
/usr/include/term.h
/usr/include/termio.h
/usr/include/time.h
/usr/include/tp_defs.h
/usr/include/tryrec.h
/usr/include/unctrl.h
/usr/include/unistd.h
/usr/include/ustat.h
/usr/include/utime.h
/usr/include/utmp.h
/usr/include/values.h
/usr/include/varargs.h
/usr/include/volhdr.h
/usr/lib/dag
/usr/lib/flip
/usr/lib/gprof.callg
/usr/lib/gprof.flat
/usr/lib/lint1
/usr/lib/lint2
/usr/lib/llib-1c
/usr/lib/llib-1c.ln
/usr/lib/llib-port
/usr/lib/llib-port.ln
/usr/lib/xcpp
/usr/lib/xpass

```

/usr/include/sys/text.h

## Fileset PFORTRAN

/system/PFORTRAN/revlist  
/usr/bin/fsplit

/usr/bin/ratfor

## Fileset PPROG

/bin/adb  
/bin/as20  
/bin/astrn  
/bin/atrans  
/bin/cdb  
/bin/fdb  
/bin/m4  
/bin/nm  
/bin/pdb  
/bin/xdb  
/etc/install  
/lib/crt0.o  
/lib/gcrt0.o  
/lib/libPW.a  
/lib/libc\_p.a  
/lib/libm.a  
/lib/libm\_p.a  
/lib/mcrt0.o  
/system/PPROG/customize  
/system/PPROG/revlist  
/usr/bin/atime  
/usr/bin/chatr  
/usr/bin/cpset  
/usr/bin/csplrit  
/usr/bin/lex  
/usr/bin/lorder

/usr/bin/mkstr  
/usr/bin/yacc  
/usr/lib/cdb.help  
/usr/lib/cdb.help.man  
/usr/lib/end.o  
/usr/lib/help/ut  
/usr/lib/lex/ncform  
/usr/lib/lex/nrform  
/usr/lib/libBSD.a  
/usr/lib/libV7.a  
/usr/lib/libcurses.a  
/usr/lib/libdbm.a  
/usr/lib/libdvio.a  
/usr/lib/libhp.a  
/usr/lib/libl.a  
/usr/lib/libmalloc.a  
/usr/lib/libndbm.a  
/usr/lib/libportnls.a  
/usr/lib/libtermcap.a  
/usr/lib/libtermlib.a  
/usr/lib/liby.a  
/usr/lib/llib-lm  
/usr/lib/llib-lm.ln  
/usr/lib/llib-lmalloc.l  
/usr/lib/xdb.help  
/usr/lib/yaccpar

## Partition MISC\_UTILS

### Fileset PACCT

/etc/newconfig/holidays  
/system/PACCT/CDFcustomize  
/system/PACCT/customize  
/system/PACCT/revlist

/usr/lib/acct/diskusg  
/usr/lib/acct/dodisk  
/usr/lib/acct/fwtmp  
/usr/lib/acct/lastlogin

/usr/lib/acct/acctcms	/usr/lib/acct/monacct
/usr/lib/acct/acctcom	/usr/lib/acct/nulladm
/usr/lib/acct/acctcon1	/usr/lib/acct/prctmp
/usr/lib/acct/acctcon2	/usr/lib/acct/prdaily
/usr/lib/acct/acctdisk	/usr/lib/acct/ptacct
/usr/lib/acct/acctdusg	/usr/lib/acct/ptecms.awk
/usr/lib/acct/acctmerg	/usr/lib/acct/ptelus.awk
/usr/lib/acct/accton	/usr/lib/acct/remove
/usr/lib/acct/acctprc1	/usr/lib/acct/runacct
/usr/lib/acct/acctprc2	/usr/lib/acct/shutacct
/usr/lib/acct/acctwtmp	/usr/lib/acct/startup
/usr/lib/acct/chargefee	/usr/lib/acct/turnacct
/usr/lib/acct/ckpacct	/usr/lib/acct/wtmpfix

## Fileset PAPPLIC

/bin/bdf	/usr/bin/factor
/bin/bs	/usr/bin/finger
/bin/chsh	/usr/bin/lastcomm
/etc/dmesg	/usr/bin/leave
/etc/dump	/usr/bin/lock
/etc/newconfig/lib.b	/usr/bin/mkmf
/etc/restore	/usr/bin/primes
/lib/libld.a	/usr/bin/printenv
/system/PAPPLIC/customize	/usr/bin/script
/system/PAPPLIC/revlist	/usr/bin/shar
/usr/bin/bc	/usr/bin/shl
/usr/bin/cal	/usr/bin/tic
/usr/bin/calendar	/usr/bin/units
/usr/bin/captainfo	/usr/bin/untic
/usr/bin/chfn	/usr/bin/uptime
/usr/bin/clear	/usr/bin/users
/usr/bin/dc	/usr/bin/vipw
/usr/bin/doschmod	/usr/bin/yes
/usr/bin/doscp	/usr/lib/calprog
/usr/bin/dosdf	/usr/lib/mf/C.L
/usr/bin/dosll	/usr/lib/mf/C.l
/usr/bin/dosln	/usr/lib/mf/C.p
/usr/bin/dosls	/usr/lib/mf/C.t
/usr/bin/dosmkdir	/usr/lib/mf/dir.p
/usr/bin/dosmv	/usr/lib/mf/man.p
/usr/bin/dosrm	/usr/lib/unittab
/usr/bin/dosrmdir	/usr/pub/ascii

## Fileset PPERIPH

/bin/mt	/system/PPERIPH/revlist
/etc/lsdev	/usr/bin/ftio

## Fileset PSCCS

/system/PSCCS/customize	/usr/bin/sact
/system/PSCCS/revlist	/usr/bin/sccsdiff
/usr/bin/admin	/usr/bin/unget
/usr/bin/cdc	/usr/bin/val
/usr/bin/ci	/usr/bin/vc
/usr/bin/co	/usr/lib/help/ad
/usr/bin/comb	/usr/lib/help/cb
/usr/bin/delta	/usr/lib/help/cm
/usr/bin/get	/usr/lib/help/cmds
/usr/bin/ident	/usr/lib/help/co
/usr/bin/merge	/usr/lib/help/de
/usr/bin/prs	/usr/lib/help/ge
/usr/bin/rcs	/usr/lib/help/prs
/usr/bin/rcsdiff	/usr/lib/help/rc
/usr/bin/rcsmmerge	/usr/lib/help/un
/usr/bin/rlog	/usr/lib/help/vc
/usr/bin/rmchg	/usr/lib/rdiff
/usr/bin/rmdel	/usr/lib/rdiff3

## Partition NLS

### Fileset NCORE

/system/NCORE/customize	/usr/bin/gencat
/system/NCORE/revlist	/usr/bin/insertmsg
/usr/bin/dumpmsg	/usr/bin/inv
/usr/bin/ex16	/usr/bin/nljust
/usr/bin/findmsg	/usr/bin/vis
/usr/bin/findstr	/usr/lib/exrecover16
/usr/bin/forder	

### Fileset NNCOMPUT

/system/NNCOMPUT/customize	/usr/lib/nls/n-computer/frecover.cat
/system/NNCOMPUT/revlist	/usr/lib/nls/n-computer/gencat.cat
/usr/lib/nls/n-computer/accept.cat	/usr/lib/nls/n-computer/getaccess.cat
/usr/lib/nls/n-computer/at.cat	/usr/lib/nls/n-computer/getopt.cat
/usr/lib/nls/n-computer/grep.cat	/usr/lib/nls/n-computer/nljust.cat

```

/usr/lib/nls/n-computer/audsys.cat /usr/lib/nls/n-computer/insertmsg.cat
/usr/lib/nls/n-computer/audusr.cat /usr/lib/nls/n-computer/ln.cat
/usr/lib/nls/n-computer/cancel.cat /usr/lib/nls/n-computer/lp.cat
/usr/lib/nls/n-computer/cat.cat /usr/lib/nls/n-computer/lpadmin.cat
/usr/lib/nls/n-computer/cdb.cat /usr/lib/nls/n-computer/lpmove.cat
/usr/lib/nls/n-computer/chacl.cat /usr/lib/nls/n-computer/lpsched.cat
/usr/lib/nls/n-computer/chgrp.cat /usr/lib/nls/n-computer/lpshut.cat
/usr/lib/nls/n-computer/chmod.cat /usr/lib/nls/n-computer/lpstat.cat
/usr/lib/nls/n-computer/chown.cat /usr/lib/nls/n-computer/ls.cat
/usr/lib/nls/n-computer/cmp.cat /usr/lib/nls/n-computer/lsacl.cat
/usr/lib/nls/n-computer/col.cat /usr/lib/nls/n-computer/mailx.cat
/usr/lib/nls/n-computer/comm.cat /usr/lib/nls/n-computer/man.cat
/usr/lib/nls/n-computer/cp.cat /usr/lib/nls/n-computer/mkdir.cat
/usr/lib/nls/n-computer/cpio.cat /usr/lib/nls/n-computer/more.cat
/usr/lib/nls/n-computer/cpp.cat /usr/lib/nls/n-computer/mv.cat
/usr/lib/nls/n-computer/cron.cat /usr/lib/nls/n-computer/neqn.cat
/usr/lib/nls/n-computer/csh.cat /usr/lib/nls/n-computer/nroff.cat
/usr/lib/nls/n-computer/cu.cat /usr/lib/nls/n-computer/od.cat
/usr/lib/nls/n-computer/date.cat /usr/lib/nls/n-computer/passwd.cat
/usr/lib/nls/n-computer/deroff.cat /usr/lib/nls/n-computer/pg.cat
/usr/lib/nls/n-computer/diff.cat /usr/lib/nls/n-computer/pr.cat
/usr/lib/nls/n-computer/diffh.cat /usr/lib/nls/n-computer/pwd.cat
/usr/lib/nls/n-computer/reject.cat /usr/lib/nls/n-computer/rm.cat
/usr/lib/nls/n-computer/ed.cat /usr/lib/nls/n-computer/rmdir.cat
/usr/lib/nls/n-computer/enable.cat /usr/lib/nls/n-computer/sccs.cat
/usr/lib/nls/n-computer/ex.cat /usr/lib/nls/n-computer/sh.cat
/usr/lib/nls/n-computer/expr.cat /usr/lib/nls/n-computer/su.cat
/usr/lib/nls/n-computer/tbl.cat /usr/lib/nls/n-computer/tee.cat
/usr/lib/nls/n-computer/uniq.cat /usr/lib/nls/n-computer/write.cat
/usr/lib/nls/n-computer/fgrep.cat /usr/lib/nls/n-computer/vis.cat
/usr/lib/nls/n-computer/find.cat /usr/lib/nls/n-computer/wc.cat
/usr/lib/nls/n-computer/forder.cat /usr/lib/nls/n-computer/xdb.cat
/usr/lib/nls/n-computer/crontab.cat
/usr/lib/nls/n-computer/audevent.cat
/usr/lib/nls/n-computer/dumpmsg.cat
/usr/lib/nls/n-computer/disable.cat
/usr/lib/nls/n-computer/expreserve.cat
/usr/lib/nls/n-computer/exrecover.cat
/usr/lib/nls/n-computer/fbackup.cat
/usr/lib/nls/n-computer/findmsg.cat
/usr/lib/nls/n-computer/findstr.cat

```

## Fileset PORT\_NLS

```
/system/PORT_NLS/APORT_NLS          /usr/lib/nls/greek/custdat.cat
/system/PORT_NLS/revlist             /usr/lib/nls/icelandic/custdat.cat
/usr/lib/nls/american/custdat.cat    /usr/lib/nls/italian/custdat.cat
/usr/lib/nls/c-french/custdat.cat    /usr/lib/nls/japanese/custdat.cat
/usr/lib/nls/chinese-s/custdat.cat   /usr/lib/nls/katakana/custdat.cat
/usr/lib/nls/chinese-t/custdat.cat   /usr/lib/nls/korean/custdat.cat
/usr/lib/nls/danish/custdat.cat      /usr/lib/nls/n-computer/custdat.cat
/usr/lib/nls/dutch/custdat.cat       /usr/lib/nls/norwegian/custdat.cat
/usr/lib/nls/english/custdat.cat     /usr/lib/nls/portuguese/custdat.cat
/usr/lib/nls/finnish/custdat.cat     /usr/lib/nls/spanish/custdat.cat
/usr/lib/nls/french/custdat.cat      /usr/lib/nls/swedish/custdat.cat
/usr/lib/nls/german/custdat.cat
```

# Contents

---

## D. Operating System Parameters

acctresume	D-2
acctsuspend	D-4
argdevnblk	D-6
dos_mem_byte	D-7
dskless_cbufs	D-8
dskless_fsbufs	D-10
dskless_mbufs	D-12
dskless_node	D-14
dst	D-15
filesizelimit	D-16
fpa	D-17
maxdsiz	D-18
maxssiz	D-19
maxswapchunks	D-20
maxtsiz	D-22
maxuprc	D-23
maxusers	D-25
mesg	D-27
minswapchunks	D-28
msgmap	D-29
msgmax	D-31
msgmnb	D-32
msgmni	D-33
msgseg	D-34
msgssz	D-36
msgtql	D-38
nbuf	D-39
ncallout	D-41

ndilbuffers . . . . .	D-43
netmeminit . . . . .	D-44
netmemmax . . . . .	D-45
netmemthresh . . . . .	D-46
nfile . . . . .	D-48
nflocks . . . . .	D-49
ngcsp . . . . .	D-50
ninode . . . . .	D-52
nproc . . . . .	D-53
npty . . . . .	D-55
ntext . . . . .	D-56
num_cnodes . . . . .	D-57
num_lan_cards . . . . .	D-59
parity_option . . . . .	D-60
scroll_lines . . . . .	D-62
selftest_period . . . . .	D-63
sema . . . . .	D-64
semaem . . . . .	D-66
semmap . . . . .	D-68
semmni . . . . .	D-70
semmns . . . . .	D-71
semmnu . . . . .	D-72
semume . . . . .	D-74
semvmx . . . . .	D-76
server_node . . . . .	D-77
serving_array_size . . . . .	D-79
shmall . . . . .	D-81
shmbrok . . . . .	D-82
shmem . . . . .	D-83
shmmax . . . . .	D-84
shmmaxaddr . . . . .	D-85
shmmin . . . . .	D-86
shmni . . . . .	D-87
shmseg . . . . .	D-88
timeslice . . . . .	D-89
timezone . . . . .	D-91
unlockable_mem . . . . .	D-92
using_array_size . . . . .	D-94



Swap Space Parameter Interaction . . . . . D-96



## Operating System Parameters

---

If you need to customize your HP-UX kernel, you can use these tunable parameters with the `config` command. `config` is described in Chapter 6, the section called “Using the config Command”.

---

**Note** If you change these parameters you can adversely affect the operation of HP-UX. *Be sure you know all the implications of using the parameters before you configure your system.*

---

**Note** Operating system parameter values may change from release to release because new functionality often requires tuning the kernel. Please be aware of this.

---

---

## **acctresume**

### **Name**

acctresume—resume accounting due to disk usage

### **Range**

-100 → 101

### **Default**

4

### **Use**

The system automatically disables process accounting when the available space on the file system where the accounting file resides falls below a certain threshold. The threshold is described under `acctsuspend`. The system also automatically re-enables process accounting when sufficient space becomes available. The parameter, `acctresume`, is the threshold (percentage of free space) which the system must have to re-enable process accounting. This percentage is added to minimum free percentage (`minfree`) for the file system.

A value of zero re-enables accounting when the free space reaches `minfree`. A value less than zero allows process accounting to use the space which is reserved for superuser use. A value greater than 100 prevents process accounting from ever being re-enabled because of available space.

When accounting is re-enabled in this way, the message:

```
"Accounting resumed"
```

is printed on the system console.

`acctresume` is only relevant to systems which turn on process accounting.

**acctresume**

**Cost**

None.

**Dependencies (interactions with other system values)**

acctsuspend < acctresume

---

## acctsuspend

### Name

acctsuspend—suspend accounting due to disk usage

### Range

-100 → 100

### Default

2

### Use

The system automatically disables process accounting when the available space on the file system where the accounting file resides falls below a certain threshold. The parameter, `acctsuspend` (specified as a percentage of free space), is the threshold. This percentage is added to minimum free percentage (`minfree`) for the file system.

A value of zero disables accounting when the free space falls below `minfree`. A value less than zero allows process accounting to use the space which is reserved for superuser use. If the sum of `acctsuspend` and `minfree` is less than zero, process accounting can never be disabled for this reason.

When accounting is disabled in this way, the message:

**"Accounting suspended"**

is printed on the system console.

`acctsuspend` is only relevant to systems which turn on process accounting.

**Cost**

None.

**Dependencies (interactions with other system values)**

`acctsuspend < acctresume`

---

## argdevnblk

### Name

argdevnblk—limit the size of the area reserved for `argdev`

### Range

0                    If 0, the area is configured dynamically.  
48 → 256            Number of disk blocks (in multiples of 12). Each disk block equals 1 Kbyte.

### Default

0

### Use

`exec` uses `argdevnblk` as a scratch pad. The bigger the `argdevnblk`, the greater the number of concurrent `execs` allowed. Each `exec` call takes approximately 12 disk blocks.

### Cost

Since `argdev` occupies the first part of the swap device, the greater the `argdevnblk`, the smaller the swap space actually available to the user's process.

### Dependencies (interactions with other system values)

None.



---

## **dos\_mem\_byte**

### **Name**

`dos_mem_byte`—Reserves memory for the HP 98286 DOS Coprocessor

### **Range**

0 → memory-limited

### **Default**

0

### **Use**

The HP 98286 DOS Coprocessor uses system memory. This memory must be reserved on Series 300 systems with more than 6 Mbytes of main memory. The reserved memory is not available to HP-UX regardless of whether the DOS Coprocessor is running.

On Series 300 with 6 Mbytes or less, reserved memory is not required. The DOS Coprocessor shares system memory with HP-UX.

The amount of reserved memory depends upon the memory configuration of the DOS Coprocessor. 1 Mbyte of reserved memory is sufficient for DOS configured with 640 Kbytes of main memory and 64 Kbytes of expanded memory. Use of more DOS expanded or DOS extended memory requires correspondingly more reserved memory.

### **Cost**

The amount of memory specified is reserved only for use by the DOS Coprocessor and is unavailable for other use.

### **Dependencies**

none

---

## dskless\_cbufs

### Name

dskless\_cbufs—number of pages allocated to the cbuf (cluster buffer) pool

### Range

0 → 128

### Default

0	If the diskless code is not configured into the kernel.
dskless_mbufs*2+1	If the diskless code is configured into the kernel. Unless you change other configurable parameters, this value will be 12.

The diskless code is configured into the kernel if the kernel description file contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

### Use

This parameter determines the number of pages allocated to the cnode's **cbuf** pool. The cbuf pool, along with the diskless mbuf pool, are used for both inbound and outbound traffic.

On the root server, the cbuf pool is allocated when the `cluster` command is issued. On a diskless cnode, the cbuf pool is allocated at kernel initialization time.

### Cost

Each cbuf is 1 Kbyte in length. There are four cbufs per page since each page on a Series 300 is 4 Kbytes. The `dskless_cbufs` parameter is actually a multiplier used as follows:

$$\text{"dskless\_cbufs"} \times \text{"pagesize"}$$

**Dependencies (interactions with other system values)**

None. The diskless cluster pool is only utilized by the diskless code. The size of the diskless cbuf pool is directly dependent upon the `dskless_mbufs` parameter ( $dskless\_mbufs*2+1$ ).

---

## **dskless\_fsbufs**

### **Name**

`dskless_fsbufs`—size of the file system buffers that can be allocated under interrupt

### **Range**

0 → 100

### **Default**

0                      If the diskless code is not configured into the kernel.

`serving_array_size`   If the diskless code is configured into the kernel.

The diskless code is configured into the kernel if the kernel description file you used contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

### **Use**

This parameter directly sizes the diskless `fsbuf` page pool. The `fsbuf` pool is a collection of file system buffers that are utilized for inbound diskless traffic. The value should be  $\leq$  `serving_array_size`. If you configure it to be  $>$  `serving_array_size`, it will be reset to be equal to `serving_array_size`. Diskless cnodes receive much less traffic than the root server, so this parameter should be smaller on the diskless cnode than on the root server.

If an inbound diskless message is large enough to require one of these buffers and none are available, the diskless protocol correctly handles retries.

On the root server the `fsbuf` pool is allocated when the root server issues the `cluster` command. On a diskless cnode, the `netbuf` pool is allocated at kernel initialization time.

**Cost**

`pagesize × dskless_fsbufs`

`pagesize` on the Series 300 is 4 Kbytes. Pages used for `fsbufs` are unavailable for general use.

**Dependencies (interactions with other system values)**

By default, `dskless_fsbufs` equals `serving_array_size`. It should always be  $\leq$  the `serving_array_size` parameter.

---

## **dskless\_mbufs**

### **Name**

`dskless_mbufs`—number pages allocated to the mbuf pool

### **Range**

0 → 64

### **Default**

0                      If the diskless code is not configured into the kernel.

`((server_node* num_cnodes)/4)+1`    If the diskless code is configured into the kernel. Unless you change other configurable parameters, this value will be 6.

Note that the diskless code is configured into the kernel if the kernel description file contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

### **Use**

This parameter determines the number of pages allocated to the diskless `mbuf`<sup>1</sup> pool. The `mbuf` pool is a collection of small buffers used for both inbound and outbound diskless traffic.

On the root server, the `mbufs` are allocated when the root server issues the `cluster` command. On a diskless `cnode`, the `mbuf` pool is allocated at kernel initialization time.

---

<sup>1</sup> The term `mbuf` is taken from the Berkeley system, but this implementation is used specifically for HP's diskless cluster, and does not directly correspond to the Berkeley `mbuf`.

## **Cost**

$\text{pagesize} \times \text{dskless\_mbufs}$

On the Series 300, `pagesize` is 4 Kbytes.

Since these pages are allocated for diskless traffic, they are unavailable for general use.

## **Dependencies (interactions with other system values)**

The diskless mbuf pool size is dependent on the parameters `server_node` and `num_cnodes`. The diskless mbuf pool is only utilized by the diskless code.

---

## **dskless\_node**

### **Name**

`dskless_node`—flag used to let system know if diskless cnode or root server.

### **Range**

1      If 1, this is a diskless cnode.

0      If 0, this is a server node.

### **Default**

0

### **Use**

If you are configuring the root server node, `dskless_node` should be 0 (default) and `server_node` should be set to 1. If you are configuring a standalone system, both `dskless_node` and `server_node` should be 0 (the default).

### **Cost**

None. There is no direct cost for this flag.

### **Dependencies (interactions with other system values)**

None. This parameter is not used; it is reserved for use in future releases.



---

## dst

### Name

dst—daylight savings time

### Range

0 → 5

### Default

1

### Use

This variable defines the daylight savings time correction:

(from file ../h/time.h)

```
#define DST_NONE      0      /* not on dst */
#define DST_USA      1      /* USA style dst */
#define DST_AUST      2      /* Australian style dst */
#define DST_WET      3      /* Western European dst */
#define DST_MET      4      /* Middle European dst */
#define DST_EET      5      /* Eastern European dst */
```

The `dst` and `timezone` parameters are used in conjunction with the `gettimeofday` system call. They do not influence the `date` command.

### Cost

None.

### Dependencies (interactions with other system values)

It is used in conjunction with `time zone`.

---

## filesizelimit

### Name

filesizelimit—sets the file size limit for process 0

### Range

0x00000010 → resource limited

### Default

0x1fffffff

### Use

This parameter sets the file size limit for process 0. The limit is given in units of 512-byte blocks and is inherited by child processes. Setting this parameter, effectively sets both the maximum and current filesize for all processes. If this parameter is not set by the user when configuring a kernel, the default setting will be 0x1fffffff, which was the value it had before the parameter was tunable. After boot-up, the file size limit can always be changed for a process and its children, using the system call, `ulimit`. (Refer to the *ulimit(2)* entry in the *HP-UX Reference*.)

### Cost

None.

### Dependencies (interactions with other system values)

None.

---

## **fpa**

### **Name**

fpa—HP 98248A Floating Point Accelerator enable/disable

### **Range**

0,1

### **Default**

1

### **Use**

fpa determines whether the code for the HP 98248A Floating Point Accelerator is included in the kernel. If `fpa = 1` the code is included, if `fpa = 0` it is not.

If you wish to use the floating point capabilities, you must:

- set `fpa` to 1 (the default)
- create a device file for the card (described in the HP 98248A installation note)
- download the floating point microcode (described in the HP 98248A installation note)

### **Cost**

Approximately 5 Kbytes.

### **Dependencies (interactions with other system values)**

None.

---

## **maxdsiz**

### **Name**

maxdsiz—maximum data size

### **Range**

262 144 bytes → 31 457 280 bytes (256 Kbytes → 30 Mbytes)

To increase beyond 30 Mbytes, see Table D-1 at the end of this Appendix.

### **Default**

0x01000000 (16 777 216 bytes, or 16 Mbytes)

### **Use**

This value is entered in bytes.

maxdsiz defines the maximum size that the data segment of an executing process can have.

The default value is large enough for the data used by most processes. maxdsiz should only be increased if you know that you have one or more processes which use huge amounts of data.

Each time the system loads a process, or an executing process attempts to expand its data segment, the system checks the size of the process's data segment. If the maxdsiz is exceeded, the process will be terminated or returned with an appropriate error.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

process text + process data + process stack ≤ 16 Mbytes (Model 310)

---

## **maxssiz**

### **Name**

maxssiz—maximum stack size

### **Range**

262 144 bytes → 31 457 280 bytes (256 Kbytes → 30 Mbytes)

To increase beyond 30 Mbytes, see table D-1 at the end of this appendix.

### **Default**

0x00200000 (2 097 152 bytes, or 2 Mbytes)

### **Use**

This value is entered in bytes.

maxssiz defines the maximum size that the stack segment of an executing process can have.

The default is large enough for the stack of most processes. maxssiz should only be increased if you have one or more processes which require a huge stack.

The stack grows dynamically. As the stack grows, the system checks the size of the process's stack segment. If the maxssiz is exceeded the process will be killed.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

process text + process data + process stack ≤ 16 Mbytes (Model 310)

---

## maxswapchunks

### Name

maxswapchunks—limit the amount of swap on a workstation

### Range

1 → ∞

### Default

512

### Use

This parameter limits the amount of swap space an individual workstation can have. This limiter should be used for a diskless cnode that tends to use a lot of the swap space connected to the root server. There is no reason to limit the amount of swap space if the diskless cnode has its own swap devices.

This limiter should be set to be larger than total swap space on a standalone system, or on a diskless cnode with local swap space. The maximum amount of swap space for the workstation is  $\text{maxswapchunks} \times \text{dmmax} \times 1024$  bytes. Using default values of `maxswapchunks` and `dmmax` gives 256 Mbytes of swap space.

It is possible to get chunks of swap space at the end of the disk that are less than `dmmax` in size. Therefore, the minimum amount of swap space that a workstation can have is:

$$(\text{maxswapchunks} - N) \times \text{dmmax}$$

where  $N$  is the number of swap devices.

**maxswapchunks**

**Cost**

$6 \times \text{maxswapchunks}$  bytes

**Dependencies**

The total swap space is limited by  $\text{maxswapchunks} \times \text{dmmmax} \times 1024$  bytes

---

## **maxtsiz**

### **Name**

`maxtsiz`—maximum text size

### **Range**

262 144 bytes → 20 971 520 bytes (256 Kbytes → 20 Mbytes)

To increase beyond 20 Mbytes, see Table D-1 at the end of this appendix.

### **Default**

0x01000000 (16 777 216 bytes, or 16 Mbytes)

### **Use**

This value is entered in bytes.

`maxtsiz` defines the maximum size that the shared text segment of an executing process can have.

The current default will accommodate the text segments of most processes. Unless you run a process with a text segment larger than 16 Megabytes, `maxtsiz` should not be modified.

Each time the system loads a process with shared text, the system checks the size of its shared text segment. `exec` issues an error message and aborts the process if the process's text segment is larger than `maxtsiz`.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

process text + process data + process stack ≤ 16 Mbytes (Model 310)



---

## maxuprc

### Name

maxuprc—maximum number of user processes

### Range

3 → (nproc - 3)

### Default

50 processes

### Use

maxuprc defines the maximum number of simultaneous processes that a user can have. A user is identified by the UID (user ID) number, not by the number of login instances. Each user will need at least one process for the shell, plus an adequate number to be able to do useful work (25 is usually more than enough).

The superuser is exempt from this limit.

Pipelines need at least one simultaneous process for each side of a '|'. Some commands, such as `cc`, `fc`, and `pc`, use more than one process per invocation.

When the total number of processes for a given user is larger than maxuprc, the following message is printed to the user tty:

```
no more processes  if from shell, or  
EAGAIN error      if a fork call from a program.
```

### Cost

none.

**maxuprc**

### **Dependencies (interactions with other system values)**

If **maxuprc** is set to a value greater than or equal to **nproc** (maximum number of processes in the system) then **maxuprc** is no longer a limit, and the system can be hoarded by a single user.

---

## maxusers

### Name

maxusers—limiter for system resource allocation

### Range

0 → memory limited

### Default

8

### Use

maxusers is a limiter for system resource allocation. The global operating system parameters that depend on maxusers are listed in the “Dependencies” section. maxusers is *not* a limiter for the number of users in the system, therefore its name (inherited from System V) is confusing.

### Cost

By itself, maxusers does not determine the size of any structures in the system. However, the value of maxusers indirectly affects kernel data structures. Consequently, the amount of memory available for user processes will be smaller with higher values of maxusers. Data structures that depend on the value of maxusers will grow as follows<sup>2</sup> :

32 bytes × nclist  
180 bytes × nproc  
16 bytes × ncallout  
286 bytes × ninode  
30 bytes × nfile  
188 bytes × ntext  
20 bytes × serving\_array\_size

<sup>2</sup> Although “maxusers” affects the “nclist” variable, “nclist” is not considered a tunable parameter. All other variables listed here are tunable and each is described elsewhere in this appendix.

## **maxusers**

### **Dependencies (interactions with other system values)**

`nclist = 100+(16×maxusers)`

`nproc = 20 + (8 × maxusers) + ngcsp`

`ncallout = 16 + nproc + using_array+size + serving_array_size`

`ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)`

`nfile = 16 × ((nproc+16+maxusers) ÷ 10) +32+ (2 × npty)`

`ntext = 40 + maxusers`

`serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)`

---

## mesg

### Name

mesg—System V messages

### Range

0 → 1

### Default

1

### Use

mesg determines whether the code for System V IPC messages is included in the kernel. If mesg=1, the code is included; if mesg=0, then the code is not included. If mesg=0, all programs using the system calls (refer to the entries for *msgget(2)*, *msgbp(2)*, and *msgctl(2)* in the *HP-UX Reference*) will receive a SIGSYS signal.

### Cost

2680 bytes.

### Dependencies (interactions with other system values)

All message parameters depend on mesg.

If mesg = 0, then the message parameters are not declared and no message structures are allocated.

If mesg = 1, then the message parameters are declared and can be tuned and all other message structures are allocated.

---

## minswapchunks

### Name

minswapchunks—set the minimum amount of swap space on a system

### Range

1 → ∞

### Default

4

### Use

This parameter is used to set the minimum amount of swap space for a system. The minimum amount of swap space would be:

$$\text{minswapchunks} \times \text{dmmax} \times 1024 \text{ bytes}$$

### Cost

There is no cost for this parameter.

### Dependencies (interactions with other system values)

Must not be greater than maxswchunks. Set to maxswchunks if larger.

Total swap space must be at least  $\text{minswapchunks} \times \text{dmmax}$

---

## msgmap

### Name

msgmap—message map

### Range

3 → memory limited

### Default

(msgtql+2)

### Use

Each set of messages allocated per identifier occupies one or more contiguous slots in the msg array. As messages are allocated and deallocated the msg array may become fragmented.

msgmap dimensions the resource map used to allocate the buffer space for messages. This map shows the free holes in the msg array. An entry in the map is used to point to each set of contiguous unallocated slots, and it consists of a pointer to the set, plus the size of the set.

If message usage is heavy, and a request for a message set cannot be accommodated, the message:

**danger: mfree map overflow**

will appear. If you get the error, you should make a new kernel with a larger msgmap.

There will be less fragmentation of the msg array if all message identifiers have the same number of messages. msgmap can then be smaller.

Three (3) is the lower limit: one slot is overhead for the map and the second slot is always needed at system initialization to indicate that the msg array is completely free.

## **msgmap**

### **Cost**

4 bytes × msgmap

### **Dependencies (interactions with other system values)**

$\text{msgmap} \leq (\text{msgtql} + 2)$

$\text{msgmap} \leq (\text{msgseg} + 2)$

If  $\text{mesg} = 0$ , then the message code is not included in the kernel and the value of  $\text{msgmap}$  is irrelevant.

If  $\text{mesg} = 1$ , then the message code is included and  $\text{msgmap}$  can then be tuned.



---

## **msgmax**

### **Name**

msgmax—message maximum size

### **Range**

0 → memory limited

### **Default**

8192 bytes

### **Use**

msgmax limits the size, in bytes, of a single message.

It should be increased only if you have applications which require larger messages. Its main value is to keep malicious or poorly written programs from using all the message buffer space.

A `msgsnd` system call which attempts to send a message larger than `msgmax` bytes returns the `EINVAL` error.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

`msgmax`  $\leq$  `msgmnb`

If `mesg` = 0, then the message code is not included in the kernel and the value of `msgmax` is irrelevant.

If `mesg` = 1, then the message code is included and `msgmax` can then be tuned.

---

## **msgmnb**

### **Name**

msgmnb—maximum number of bytes on message queue

### **Range**

0 → memory limited

### **Default**

16 384 bytes

### **Use**

msgmnb is the maximum total size, in bytes, of all messages that can be queued on a message queue at the same time.

A msgsnd system call which attempts to exceed this limit either:

- blocks or,
- returns EAGAIN error

depending on whether the IPC\_NOWAIT flag is set with the call.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

msgmnb  $\geq$  msgmax

msgmnb  $\leq$  (msgssz $\times$ msgseg)

If msg = 0, then the message code is not included in the kernel and the value of msgmnb is irrelevant.

If msg = 1, then the message code is included and msgmnb can then be tuned.

---

# msgmni

## Name

msgmni—number of message queue identifiers

## Range

1 → memory limited

## Default

50

## Use

msgmni dimensions the array of message queue identifiers.

A message queue identifier is needed for each message queue in the system.

When msgmni message queues already exist, an attempt to allocate a new message queue with the `msgget` system call returns an ENOSPC error.

If a message queue is not deallocated, it will continue to exist in the system even after the process(es) using it have died.

When you are through using a message queue, remember to deallocate them. Running into the msgmni limit usually means that users have not freed them up.

## Cost

46 bytes×msgmni

## Dependencies (interactions with other system values)

If `mesg = 0`, then the message code is not included in the kernel and the value of msgmni is irrelevant.

If `mesg = 1`, then the message code is included and msgmni can then be tuned.

---

## **msgseg**

### **Name**

`msgseg`—message segments

### **Range**

1 → 32 767

### **Default**

16 384

### **Use**

`msgseg` together with `msgssz` determines the size of the buffer available for queuing messages.

`msgssz` determines the size, in bytes, of the units in which messages are allocated space. When a message is allocated, its size is rounded up to the nearest multiple of `msgssz`.

`msgseg` is the number of these units available.

In most cases the product of `msgseg`×`msgssz` is of most interest since it determines the total amount of space available for messages. Different `msgseg:msgssz` ratios which yield the same product will just cause this space to be fragmented differently for the same usage.

### **Cost**

`msgssz` bytes×`msgseg`

**Dependencies (interactions with other system values)**

$(\text{msgseg} \times \text{msgssz}) \geq \text{msgmnb}$

$(\text{msgseg} \times \text{msgssz}) \geq \text{msgmax}$

If  $\text{msg} = 0$ , then the message code is not included in the kernel and the value of  $\text{msgseg}$  is irrelevant.

If  $\text{msg} = 1$ , then the message code is included and  $\text{msgseg}$  can then be tuned.

---

## **msgssz**

### **Name**

msgssz—message segment size

### **Range**

1 → memory limited

### **Default**

1 byte

### **Use**

msgssz together with msgseg determines the size of the buffer available for queueing messages.

msgssz determines the size, in bytes, of the units in which messages are allocated space. When a message is allocated, its size is rounded up to the nearest multiple of msgssz.

msgseg is the number of these units available.

In most cases, the product of  $\text{msgseg} \times \text{msgssz}$  is of most interest since it determines the total amount of space available for messages. Different msgseg:msgssz ratios which yield the same product will just cause this space to be fragmented differently for the same usage.

### **Cost**

$\text{msgseg} \times \text{msgssz}$  bytes

**Dependencies (interactions with other system values)**

$(\text{msgseg} \times \text{msgssz}) \geq \text{msgmnb}$

$(\text{msgseg} \times \text{msgssz}) \geq \text{msgmax}$

If  $\text{mesg} = 0$ , then the message code is not included in the kernel and the value of  $\text{msgssz}$  is irrelevant.

If  $\text{mesg} = 1$ , then the message code is included and  $\text{msgssz}$  can then be tuned.

---

## **msgtql**

### **Name**

`msgtql`—number of message headers

### **Range**

1 → memory limited

### **Default**

40

### **Use**

`msgtql` dimensions an array of message headers. A message header is used for each message queued in the system.

A `msgsnd` system call which attempts to exceed this limit either:

- blocks or,
- returns EAGAIN error

depending on whether the `IPC_NOWAIT` flag is set with the call.

### **Cost**

12 bytes×`msgtql`

### **Dependencies (interactions with other system values)**

`msgmap` ≤ `msgtql` + 2

If `mesg` = 0, then the message code is not included in the kernel and the value of `msgtql` is irrelevant.

If `mesg` = 1, then the message code is included and `msgtql` can then be tuned.



---

## **nbuf**

### **Name**

nbuf—number of buffer headers

### **Range**

- 0                    At boot time, if `nbuf = 0`, the amount of buffer space is dynamically set to 10 percent of available memory for the first 5 Mbytes, 5 percent for the remaining memory.
- 16 → 768           The number of file system buffer headers is set to the specified value.

### **Default**

0 (configured dynamically)

### **Use**

The file system buffer cache is a pool of pages (`bufpages`) and a configurable number of buffer headers. `nbuf` defines the number of file system buffer headers.

`nbuf` can be left alone, in which case the default values will be used. However, `nbuf` can be used to decrease or increase the buffer cache within the ranges given above (refer to `Range`). These two cases are discussed below:

If you have a Model 310, it assigns (in pages) to the system buffer cache 10% of memory for the first 2 Mbytes. If the amount of RAM is greater than 2 Mbytes, use an `nbuf` of 51, which is the default for 2 Mbytes.

- The default case:

At boot time, as the system dynamically allocates memory, it assigns (in pages) to the system buffer cache 10 percent of the available memory for the first 5 Mbytes, 5 percent for the remaining memory<sup>1</sup>. Then, `nbuf` is computed so that each `nbuf` header has one 4 Kbyte (4096 bytes) `bufpages` of space assigned to it.

## **nbuf**

For example, on a Series 300 with 8 Mbytes of RAM, the default is `nbuf = 0`. After the dynamic allocation, `nbuf = 166` and the buffer cache =  $166 \times 4096$  bytes.

- To decrease or increase the buffer cache:

The user can specify the size of the buffer cache, independent of memory size. For instance, to specify 127 buffer headers, the sample `dfile` entry would be:

```
* tunable system parameters
  nbuf      127
```

Notice that in the above example, the size (127) may be smaller or larger than the default case.

These buffers are used for all file system I/O operations, plus all other block I/O operations in the system (for example, `exec`, `mount`, i-node reading, some device drivers, etc.).

## **Cost**

4096 bytes×each bufpage

104 bytes×each nbuf header structure

104 bytes×max(256, (nbuf÷2))

## **Dependencies (interactions with other system values)**

None.

---

## **ncallout**

### **Name**

`ncallout`—number of callouts

### **Range**

6 → memory limited

### **Default**

$16 + \text{nproc} + \text{using\_array\_size} + \text{serving\_array\_size}$

### **Use**

`ncallout` is the maximum number of timeouts which can be scheduled by the kernel at any one time. Timeouts are used by :

- “alarm” (system call)
- `setitimer` (system call)
- drivers
- `uucp` processes
- process scheduling
- HP-UX cluster kernel code

When the system runs out of callouts, it prints the following fatal error to the console:

```
panic: timeout table overflow
```

### **Cost**

16 bytes × `ncallout`

**ncallout**

## **Dependencies (interactions with other system values)**

The larger nproc is, the larger ncallout should be.

A rough guideline of 1 callout per process should be used unless one has processes which use lots of callouts.

---

## **ndilbuffers**

### **Name**

ndilbuffers—number of DIL buffers

### **Range**

1 → memory-limited

### **Default**

30

### **Use**

ndilbuffers defines the maximum number of DIL open device files at any one time in the system.

ndilbuffers is used exclusively by the Device I/O Library. If DIL is not used, no DIL buffers are necessary.

### **Cost**

314 bytes×ndilbuffers

### **Dependencies (interactions with other system values)**

None.

---

## netmeminit

### Name

`netmeminit`—number of bytes of memory to be preallocated at system initialization time

### Range

0 → 512 × NETCLBYTES (1000 Kbytes)

### Default

0

### Use

The `netmeminit` parameter specifies the number of bytes of memory to be preallocated for networking at system initialization time. There is no delay associated with networking getting physical memory from the system.

### Dependencies (interactions with other system values)

`netmeminit` < `netmemmax`

---

## **netmemmax**

### **Name**

netmemmax—network memory pool maximum

### **Range**

75 000 → 1 000 000 bytes

### **Default**

250 000 bytes

### **Use**

`netmemmax` defines the maximum network memory pool size. The network memory pool starts with minimal memory and grows as the amount of network activity grows, until `netmemmax` is reached. If more network memory is needed after the limit is reached, the network operation requesting the memory receives an “out-of-memory” message and the operation fails.

### **Cost**

The network memory pool is allocated from the available system memory. Larger values of `netmemmax` allow the network to acquire more memory, leaving less memory available for other system activities.

### **Dependencies (interactions with other system values)**

If `netmemmax` is less than or equal to `netmemthresh`, network memory reservation is always in effect.

---

## netmemthresh

### Name

netmemthresh—network memory reservation threshold

### Range

-1 → 2 147 483 647

### Default

100 000

### Use

netmemthresh defines how large the network memory pool is allowed to grow before disabling network memory reservation.

As long as network memory reservation is in effect (i.e., netmemthresh has not been reached), each connection is guaranteed its maximum memory requirements. When the cumulative effect of each connection's memory reservations meets or exceeds netmemthresh, network memory reservation is “disabled”. Subsequently, all connections are relegated to memory allocation on an “as needed” basis. Because connections are no longer guaranteed memory allocation, connections could be lost due to insufficient memory.

Network memory reservation can be enforced regardless of the size of the network memory pool by assigning netmemthresh the value -1.

### Cost

None.



**Dependencies (interactions with other system values)**

If `netmemthresh` is less than `netmemmax`, the network memory pool typically grows very quickly to `netmemthresh` size and remains there, except under extremely heavy network loads. Increasing the value of `netmemthresh` reserves more memory under light network loads, leaving less memory available to other system activities.

If `netmemthresh` is greater than or equal to `netmemmax`, or `netmemthresh` equals `-1`, network memory reservation is in effect at all times. In this case, the network memory pool typically grows to `netmemmax` size very quickly. Memory requirements of all connections are guaranteed, but fewer connections can be supported than with lower values of `netmemthresh`.

---

## nfile

### Name

nfile—number of files

### Range

14 → memory-limited

### Default

$16 \times ((nproc+16+maxusers) \div 10) + 32 + (2 \times npty)$

### Use

nfile defines the maximum number of open files at any one time in the system.

It is the number of slots in the file descriptor table. Be generous with this number since the cost is low, and not having enough slots would cut down in the amount of work that can be done simultaneously in the system.

### Cost

30 bytes×nfile

### Dependencies (interactions with other system values)

Windows/9000      3 or 4 file descriptors per window

Processes          At least 3 file descriptors per process (stdin, stdout, stderr)

Pipes              2 per pipe (1 per side)

Shell scripts and background processes are treated like any other processes.

nfile depends on nproc, maxusers, and npty as shown in the “Default” section.

---

## **nflocks**

### **Name**

nflocks—number of file locks

### **Range**

2 → 2000

### **Default**

200

### **Use**

nflocks gives the possible number of file/record locks in the system. When choosing this number it should be noted that one file may have several locks and that especially databases may need a large number of locks (if they use lockf at all).

### **Cost**

22 bytes×nflocks

### **Dependencies (interactions with other system values)**

None.

---

## ngcsp

### Name

ngcsp—number of general cluster server processes (for diskless server)

### Range

0 → memory limited

### Default

$8 \times \text{num\_cnodes}$

The default for `num_cnodes` is 0, so unless you have changed `num_cnodes` the value for `ngcsp` is also 0.

### Use

`ngcsp` specifies the maximum number of general cluster server processes (GCSPs) that can exist simultaneously in the system. GCSPs are required only on the cluster's root server, so the value of `ngcsp` should be zero on machines that run standalone or diskless.

GCSPs are started with the `csp` command; the command returns an error if asked to start more than `ngcsp` processes. The system will sometimes attempt to spawn more GCSPs than requested by this command to keep up with heavy demands. Thus it is recommended, but not required, that the value of `ngcsp` be about twice the value specified for the `csp` command.

### Cost

18 bytes  $\times$  `ngcsp`

**Dependencies (interactions with other system values)**
$$\text{ngcsp} = 8 \times \text{num\_cnodes}$$
$$\text{nproc} = 20 + 8 \times \text{maxusers} + \text{ngcsp}$$

Each GCSP requires one process slot. Therefore, **ngcsp** should be less than  $(\text{nproc} - 5)$ .

Each GCSP in use requires a serving array entry as do some other requests, therefore **ngcsp** must be less than **serving\_array\_size**.

---

## ninode

### Name

ninode—number of inodes

### Range

14 → memory-limited

### Default

$((nproc+16+maxusers)+32+(2 \times npty)+server\_node \times 18 \times num\_cnodes)$

### Use

ninode defines the maximum number of open inodes which can be in-core.

ninode specifies the maximum number of slots in the in-core inode table. This in-core table is used for fast access of inode information. For efficiency reasons, the table is hashed.

### Cost

286 bytes  $\times$  ninode

### Dependencies (interactions with other system values):

Each unique open file has an open inode associated with it, therefore the larger the number of unique open files, the larger ninode should be. If this is the root server of an HP-UX cluster this pertains to all opens on the entire cluster.

The ninode parameter depends on nproc, maxusers, npty, server\_node and num\_cnodes as shown in the equation in the “default” section.

---

## nproc

### Name

nproc—number of processes

### Range

6 → memory limited

### Default

$20+8 \times \text{maxusers} + \text{ngcsp}$

maxusers defaults to 8 and ngcsp defaults to 0, so this value defaults to 84.

### Use

nproc specifies the maximum total number of processes that can exist simultaneously in the system.

There have to be at least three (3) system overhead processes at all times (swapper, init, and pageout demon), and there is always one entry reserved for the superuser.

Examples of other features and the number of processes are:

LAN network	3 processes
Printer scheduler	1 process
Cron	1 process

Your application's manual may discuss the number of processes used. For a discussion of the number of processes used by HP Windows/9000 refer to the *HP Windows/9000 User's Manual*, the appendix "Window Limitations".

When the total number of processes in the system is larger than nproc, the following messages are printed:

- To the console:

```
proc: table is full
```

## nproc

- To the tty of the user who tries to start the last process(es):

no more processes    *if from shell, or*  
EAGAIN error        *if a fork call from a program.*

## Cost

180 bytes × nproc

## Dependencies (interactions with other system values)

ninode = (nproc + 16 + maxusers) + 32 + (2 × npty) + (server\_node × 18 × num\_cnodes)

nfile = 16 × ((nproc + 16 + maxusers) ÷ 10) + 32 + (2 × npty)

ncallout = 16 + nproc + using\_array\_size + serving\_array\_size

maxuprc ≤ (nproc - 3)

nproc = 20 + (8 × maxusers) + ngcsp

using\_array\_size is directly dependent upon the nproc configurable parameter



---

## **npty**

### **Name**

npty—number of pseudo-teletypes

### **Range**

1 → memory limited

### **Default**

82

### **Use**

npty controls the maximum number of pseudo-teletype (pty) devices in the system.

npty is used primarily by the window subsystem. If you use windows, you are a large user of npty. If you use `remsh`, `vt`, or `rlogin`, you are a small user of npty. If windows are not used, and the pty driver is not explicitly used, no pty structures are necessary.

### **Cost**

188 bytes×npty

### **Dependencies (interactions with other system values)**

None.

---

## **n<sub>text</sub>**

### **Name**

`ntext`—number of shared text descriptors

### **Range**

10 → memory limited

### **Default**

40 + `maxusers`

### **Use**

`ntext` defines the maximum number of shared text (or code) descriptors (data structures describing the text) which can be active at any one time. Note that this does not limit the number of processes sharing the same shared text program.

Attempting to start a new process that may require a new text descriptor when no more text descriptors are available will get “text table is full” message on the console and the new process will be killed.

### **Cost**

192 bytes × `ntext`

### **Dependencies (interactions with other system values)**

The default for `ntext` is 40 + `maxusers`.

---

## num\_cnodes

### Name

num\_cnodes—limiter for diskless system resource allocation

### Range

0 → 255

### Default

0

### Use

num\_cnodes is used as a limiter for diskless resource allocation; the value of other global system parameters depends on num\_cnodes. By itself it does not determine the size of any structures in the system.

It is used as an indicator of the number of diskless cnodes that a server can reasonably expect to serve simultaneously. It is *not* used as a limit for the number of cnodes supported by the server.

Suggested values are:

```
5 if 0- 5 cnodes
10 if 6-10 cnodes
15 if 11-15 cnodes
20 if 16-20 cnodes
```

### Cost

There is no direct cost for this flag.

A higher value of num\_cnodes indirectly results in larger kernel data structures and buffers; consequently, the amount of memory available for user processes becomes smaller.

**num\_cnodes**

## **Dependencies (interactions with other system values)**

`serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)`

`ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)`

`ngcsp = 8 × num_cnodes`

`dskless_mbufs = (((server_node × num_cnodes)/4)+1)`

---

## **num\_lan\_cards**

### **Name**

num\_lan\_cards—maximum number of LAN interface cards the system will support

### **Range**

0 → 5

### **Default**

2

### **Use**

num\_lan\_cards defines the maximum number of LAN interface cards the system will support. This number should be greater than or equal to the number of LAN interface cards actually present.

### **Cost**

Approximately 1200 bytes per LAN interface card.

### **Dependencies (interactions with other system values)**

None.

---

## parity\_option

### Name

parity\_option—used to handle parity errors

### Range

0 → 2

### Default

2

### Use

parity\_option selects the kind of action that the system takes if it encounters a parity error.

The actions are as follows:

- 0        Print a 'Parity error' message to the console.
- 1        Print a 'Parity error' message to console, plus:
  - if user state, it kills the current process (which may not always be the process which caused the error, as with a DMA card) and prints an error message to its tty.
  - if supervisor state, it panics with a 'parity error' message to the console.
- 2        Always panics with a 'parity error' message to console.

---

**Caution**        Values other than 2 could result in data corruption depending on where the RAM parity error occurs.

---

**Cost**

None.

**Dependencies (interactions with other system values):**

None.

---

## **scroll\_lines**

### **Name**

`scroll_lines`—total number of ITE text lines

### **Range**

100 → 5000

### **Default**

100

### **Use**

The ITE (Internal Terminal Emulator) will remember more lines of text than can be displayed on the screen. This parameter controls how many lines of text will be remembered.

For example, on a 46 line screen, if `scroll_lines` is 460, then there will be a total of ten screens of information available, of which only one is displayed at a time.

### **Cost**

$2 \times \langle \text{screen width in characters} \rangle \times \text{scroll\_lines}$

### **Dependencies (interactions with other system values)**

None.



---

## **selftest\_period**

### **Name**

`selftest_period`—period between execution of kernel selftest

### **Range**

0, 90 → 3600 seconds

### **Default**

0                    If the diskless code is not configured into the kernel.

120 sec.            If the diskless code is configured into the kernel.

Note that the diskless code is configured into the kernel if the kernel description file contains the driver called `diskless`. Refer to the section on kernel configuration for more information.

### **Use**

The diskless kernel's selftest code uses the `selftest_period` parameter to determine how often to run a selftest. The selftest checks the availability of kernel resources required for diskless operation. The parameter is interpreted as the minimum number of seconds between execution of the selftest check. A 0 value indicates that the selftest should not be run.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

None.

---

## **sema**

### **Name**

sema —System V semaphores

### **Range**

0 → 1

### **Default**

1

### **Use**

sema determines whether the code for System V IPC semaphores will be included in the kernel.

If sema = 1, the code is included;  
if sema = 0, then the code is not included.

HP Windows/9000 and the Starbase graphics library both require the semaphore code.

If sema=0, and you have programs which use the *semget(2)* or *semop(2)* system calls, you will receive a SIGSYS signal.

### **Cost**

4240 bytes.

### **Dependencies (interactions with other system values)**

All semaphore parameters depend on the value of sema.

If sema = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

## **sema**

If **sema** = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

---

**Note**            If **sema** = 0, then HP Windows/9000 and Starbase graphics will be unable to run.

---

---

## **semaem**

### **Name**

semaem—“adjust on exit” maximum value for semaphores

### **Range**

0 → min (semvmx, 32767)

### **Default**

16 384

### **Use**

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

semaem is the maximum value by which a semaphore can be undone.

This value is cumulative per process, so if one process has more than one undo operation on a semaphore, the value of each undo operation is added up in the variable semadj. semadj is the number by which the semaphore will be incremented or decremented if the process dies.

Refer to the entry for *semop(2)* in the *HP-UX Reference* for more detailed information on semaphore undos.

Any semop calls which attempt to set  $|\text{semadj}| > \text{semaem}$  result in an ERANGE error.

### **Cost**

None.

**Dependencies (interactions with other system values)**

$|\text{semaem}| \leq \text{semvmx}$

If  $\text{sema} = 0$ , then the semaphore parameters are not declared and no semaphore structures allocated.

If  $\text{sema} = 1$ , then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

---

## semmap

### Name

semmap—semaphore map

### Range

4 → memory limited

### Default

semnmi + 2

### Use

Each set of semaphores allocated per identifier occupies one or more contiguous slots in the `sem` array. As semaphores are allocated and deallocated the `sem` array may become fragmented.

`semmap` dimensions the resource map which shows the free holes in the `sem` array. An entry in this map is used to point to each set of contiguous unallocated slots, and it consists of a pointer to the set, plus the size of the set.

If semaphore usage is heavy, and a request for a semaphore set cannot be accommodated, the message:

**danger: mfree map overflow**

will appear. It will then be helpful to make a new kernel with a larger `semmap`.

Fragmentation of the `sem` array is reduced if in usage of the system all semaphore identifiers have the same number of semaphores. `semmap` can then be somewhat smaller.

Four (4) is the lower limit: 1 slot is overhead for the map and the second slot is always needed at system initialization to indicate that the `sem` array is completely free.

**Cost**

4 bytes × semmap

**Dependencies (interactions with other system values)**

(semmap-2) = maximum number of contiguous unallocated pieces of the sem array.

semmap ≤ (semnmi-2)

If sema = 0, then the semaphore code is not included in the kernel and the value of semmap is irrelevant.

If sema = 1, then the semaphore code is included and semmap can then be tuned.

---

## **semnmi**

### **Name**

`semnmi`—semaphore number of identifiers

### **Range**

2 → memory limited

### **Default**

64

### **Use**

`semnmi` defines the number of sets (identifiers) of semaphores available to the users.

When the system runs out of semaphore sets, the `semget` system call will return a `ENOSPC` error message.

### **Cost**

30 bytes × `semnmi`

### **Dependencies (interactions with other system values)**

`semnmi` ≤ `semnms`

`semnms` ≤ (`semnmi` × `semmsl`)

`semmsl` = 50 /\* maximum number of semaphores per ID, not tunable \*/

`semmap` ≤ (`semnmi` + 2)

If `sema` = 0, then the semaphore code is not included in the kernel and the value of `semnmi` is irrelevant.

If `sema` = 1, then the semaphore code is included and `semnmi` can then be tuned.



---

## **semms**

### **Name**

`semms`—semaphore number in system

### **Range**

2 → memory limited

### **Default**

128

### **Use**

`semms` defines the total number of semaphores available to the users of the system.

When the system does not have enough contiguous semaphores in the sem array to satisfy a `semget` request, the call returns a `ENOSPC` error. This error may occur even though there may be enough free semaphores, but they are not contiguous.

### **Cost**

8 bytes × `semms`

### **Dependencies (interactions with other system values)**

`semms` ≤ `semms`

`semms` ≤ (`semms` × `semmsl`)

`semmsl` = 50 /\* maximum number of semaphores per ID, not tunable \*/

If `sema` = 0, the semaphore code is not included in the kernel and the value of `semms` is irrelevant.

If `sema` = 1, the semaphore code is included and then `semms` can be tuned.

---

## semnmu

### Name

semnmu—semaphore number of undo structures

### Range

1 → (nproc - 3)

### Default

30

### Use

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

semnmu is the number of processes which can have undo's pending on a given semaphore. It determines the size of the `sem_undo` structure.

Refer to the entry for *semop(2)* for a more detailed explanation of undos.

You should increase `semume` if the user gets ENOSPC errors on `semop` calls using the SEM\_UNDO flag.

### Cost

$(6 + (8 \times \text{semume})) \text{ bytes} \times \text{semnmu}$

### Dependencies (interactions with other system values)

semnmu determines the size of the structure `sem_undo`, which in turn contains the substructure dimensioned by `semume`.

There is no point in having `semnmu > (nproc-3)` since it is the largest number of processes in the system which could use semaphores simultaneously.

## **semmnu**

If `sema = 0`, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema = 1`, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

---

## semume

### Name

semume—semaphore undo entries per process

### Range

1 → semmns

### Default

10

### Use

An undo is a special, optional, flag in a semaphore operation which causes that operation to be undone if the process which invoked it dies.

semume limits the number of semaphores that each process can have undos pending on.

Refer to the entry for *semop(2)* for a more detailed explanation of undos.

When the user gets EINVAL errors on semop calls with the SEM\_UNDO flag, then you should increase the value of semume.

### Cost

$(6+(8 \times \text{semume}))$  bytes  $\times$  semmnu

### Dependencies (interactions with other system values)

semume  $\leq$  semmns

semume is the size of the substructure undo, which is part of the sem\_undo structure. The size of sem\_undo is determined by semume.

## **semume**

If `sema = 0`, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema = 1`, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

---

## **semvmx**

### **Name**

`semvmx`—semaphore maximum value

### **Range**

1 → 65 535

### **Default**

32 767

### **Use**

`semvmx` is the maximum value that a semaphore is allowed to reach. This limit is imposed by the largest number that can be stored in a 16-bit unsigned integer (65 535).

`semop` system calls which try to increment a semaphore's value > `semvmx` result in `ERANGE` errors. If `semvmx` > 65 535 then semaphore values can overflow and these errors will not be caught.

### **Cost**

None.

### **Dependencies (interactions with other system values):**

`semaem` <= `semvmx`

If `sema` = 0, then the semaphore parameters are not declared and no semaphore structures allocated.

If `sema` = 1, then the semaphore parameters are declared and can be tuned and the semaphore structures allocated.

---

## **server\_node**

### **Name**

`server_node`—flag used to help size an array for the root server's inbound requests.

### **Range**

- 1 If 1, `serving_array[]` and `ninode` are sized for a server node.
- 0 If 0, `serving_array[]` and `ninode` are sized for a diskless cnode.

### **Default**

0

### **Use**

This parameter is used as a flag to determine the size of `serving_array[]` and `ninode`. `serving_array[]` is an array of kernel structures used by a cnode for inbound requests. If this parameter is set (`server_node=1`), `serving_array[]` will be sized appropriately for a root server node. The cluster's root server should have this parameter set to 1, and should have the `dskless_node` parameter set to 0 (default).

If you are configuring a diskless cnode, `server_node` should be 0 (default) and `dskless_node` should be set to 1. If you are configuring a standalone system, both `dskless_node` and `server_node` should be 0 (the default).

Refer to the `serving_array_size` and `ninode` operating system parameters for information on these resources.

**server\_node**

## **Cost**

None. There is not direct cost for this flag.

## **Dependencies (interactions with other system values)**

Used in `ninode` and `serving_array_size` as follows:

```
serving_array_size = (server_node × num_cnodes × maxusers + 2 × maxusers)
```

```
ninode = (nproc + 16 + maxusers)+32+(2 × npty)+(server_node × 18 × num_cnodes)
```

```
dskless_mbufs = (((server_node × num_cnodes)/4)+1)
```



---

## servicing\_array\_size

### Name

`servicing_array_size`—size of the cluster's servicing array

### Range

0 → 200

### Default

0	If the diskless code is not configured into the kernel.
$(\text{server\_node} \times \text{num\_cnodes} \times \text{maxusers}) + (2 \times \text{maxusers})$	If the diskless code is configured into the kernel. Unless you change other configurable parameters, this value will be 16.

The diskless code is configured into the kernel if the kernel description file you used contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

### Use

`servicing_array_size` defines the size of the kernel's servicing array. `servicing_array[]` is an array of kernel structures that holds information related to inbound requests. It is sized according to the following formula:

$$\text{servicing\_array\_size} = (\text{server\_node} \times \text{num\_cnodes} \times \text{maxusers} + 2 \times \text{maxusers})$$

Each inbound request consumes a single `servicing_array[]` entry that is held until the request is satisfied.

If both configurable parameters, `dskless_node` and `server_node`, are equal to zero, the system is treated as a standalone, non-clustered, kernel. In this situation, `servicing_array_size` is zero and `servicing_array[]` is not compiled into the kernel.

## **serving\_array\_size**

### **Cost**

20 bytes × `serving_array_size`

### **Dependencies (interactions with other system values)**

The size of the serving array is directly dependent upon the following parameters: `server_node`, `num_cnodes`, and `maxusers`.

$$\text{serving\_array\_size} = (\text{server\_node} \times \text{num\_cnodes} \times \text{maxusers} + 2 \times \text{maxusers})$$

By default, the number of file system buffers that can be allocated under interrupt is equal to the value of `serving_array_size` (see configurable parameter `dskless_fsbufts`).

Each GCSP requires a serving array entry (refer to `ngcsp`)

---

## **shmall**

### **Name**

shmall—total size of all shared memory segments

shmall is present just for compatibility reasons with System V.

---

## shmbrik

### Name

shmbrik—shared memory break

### Range

0 → maximum number of pages in the data segment.

### Default

16 pages

### Use

shmbrik defines the size of the gap, in pages, between the top of the current data segment and the first default shared memory address.

### Cost

None.

### Dependencies (interactions with other system values)

If `shmem = 0`, then the code for shared memory is not included in the kernel and shmbrik has no effect.

If `shmem = 1`, then the code for shared memory is included in the kernel and shmbrik is tunable.

---

## **shmem**

### **Name**

shmem—System V shared memory enable/disable

### **Range**

0, 1

### **Default**

1

### **Use**

shmem determines whether the code for System V IPC shared memory is included in the kernel. If `shmem = 1` then the code is included, if `shmem = 0` it is not.

### **Cost**

Approximately 11 Kbytes.

### **Dependencies (interactions with other system values)**

HP Windows/9000 and the Starbase graphics library both require the presence of shared memory code.

---

## **shmmax**

### **Name**

shmmax—shared memory maximum

### **Range**

2 Mbytes → 6 Mbytes

To increase beyond 6 Mbytes, see Table D-1 at the end of this appendix.

### **Default**

0x00600000 (6 Mbytes)

### **Use**

shmmax defines the maximum shared memory segment size.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

$shmmin \leq shmmax$

If  $shmем = 0$ , then the code for shared memory is not included in the kernel and the value of  $shmmax$  is irrelevant.

If  $shmем = 1$ , then the code for shared memory is included and  $shmmax$  is tunable.

For Windows/9000 to run,  $shmmax$  must be  $\geq 2$  Mbytes.

---

## **shmmaxaddr**

### **Name**

`shmmaxaddr`—shared memory maximum address

### **Range**

positive integer

### **Default**

0x01000000 (16 Mbytes)

### **Use**

`shmmaxaddr` specifies the highest address allowed for a shared memory segment in user address space.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

If `shmem = 0`, then the code for shared memory is not included in the kernel and `shmmaxaddr` has no effect.

If `shmem = 1`, then the code for shared memory is included and `shmmaxaddr` is tunable.

If you allow a shared memory segment to attach at a high address location, you will require more system overhead for the translation table data structure. *Do not reconfigure this unless absolutely necessary.*

---

## **shmmin**

### **Name**

shmmin—shared memory minimum

### **Range**

all positive integers

### **Default**

1 byte

### **Use**

shmmin defines the minimum shared memory segment size.

### **Cost**

None.

### **Dependencies (interactions with other system values)**

`shmmin < shmmax`

If `shmем = 0`, then the code for shared memory is not included in the kernel and the value of `shmmin` is irrelevant.

If `shmем = 1`, then the code for shared memory is included and `shmmin` is tunable.

If it is reconfigured other subsystems (such as Windows/9000) may not work.



---

## **shmmni**

### **Name**

shmmni—shared memory maximum number of identifiers

### **Range**

positive integers

### **Default**

30 identifiers

### **Use**

shmmni defines the system-wide maximum number of shared memory segments.

shmmni should be large enough to hold as many shared memory segments as will be used simultaneously.

### **Cost**

The data structure associated at each shared memory segment is about 100 bytes.

### **Dependencies (interactions with other system values)**

If `shmem = 0`, then the code for shared memory is not included in the kernel and `shmmni` has no effect.

If `shmem = 1`, then the code for shared memory is included and `shmmni` can be tuned.

Windows/9000 and Starbase graphics require `shmmni` to be  $\geq 4$ .

---

## **shmseg**

### **Name**

shmseg—shared memory segments

### **Range**

all positive integers

### **Default**

10

### **Use**

shmseg defines the maximum number of shared memory segments that can be attached to a process at any given time.

### **Cost**

about  $12 \text{ bytes} \times \text{nproc} \times \text{shmseg}$

### **Dependencies (interactions with other system values)**

If `shmem = 0`, then the code for shared memory is not included in the kernel and `shmseg` has no effect.

If `shmem = 1`, then the code for shared memory is included and `shmseg` can be tuned.

Windows/9000 and Starbase graphics require that `shmseg` be  $\geq 4$ .

---

## **timeslice**

### **Name**

timeslice—scheduling timeslice interval

### **Range**

-1 → 2 147 483 647

### **Default**

0

### **Use**

The system performs round-robin scheduling among processes at a given priority. The timeslice interval is the amount of time one process is allowed to run before the CPU is given to the next process at the same priority.

The value of timeslice is specified in units of 20 millisecond clock ticks. There are two special values:

- 0        use the system default value (currently 5 ticks, or 100 milliseconds)
- 1        disable round-robin scheduling completely

One side-effect of this parameter is that a process always checks for pending signals when its timeslice expires. This guarantees that a process which does not make any system calls (including a runaway process in an infinite loop) can be terminated. Thus setting timeslice to a very large value, or to -1, can prevent such processes from seeing signals.

It is anticipated that this parameter will only be changed on systems dedicated to applications with specific real-time needs.

## **timeslice**

### **Cost**

There is no memory allocation related to this parameter. There is some amount of CPU time spent at each timeslice interval, but this time has not been precisely measured.

### **Dependencies (interactions with other system values)**

None.

---

## timezone

### Name

timezone—minutes west of Greenwich

### Range

-780 → 780

### Default

420

### Use

Time zone information indicates the minutes west of Greenwich.

```
struct timezone tz = { TIMEZONE, DST };
struct timezone {
int tz_minuteswest; /* minutes west of Greenwich */
int tz_dsttime;    /* type of dst correction */
};
#define DST_NONE      0    /* not on dst */
#define DST_USA       1    /* USA style dst */
#define DST_AUST      2    /* Australian style dst */
#define DST_WET       3    /* Western European dst */
#define DST_MET       4    /* Middle European dst */
#define DST_EET       5    /* Eastern European dst */
```

The `dst` and `timezone` parameters are used in conjunction with the `gettimeofday` system call. They do not influence the `date` command.

### Cost

None.

### Dependencies (interactions with other system values)

It is used in conjunction with `dst` (daylight savings time).

---

## unlockable\_mem

### Name

unlockable\_mem—unlockable memory

### Range

0 → (the available memory indicated at power-up)

### Default

102 400 bytes

### Use

unlockable\_mem defines the minimum amount of memory which is guaranteed to be available for virtual memory and/or system overhead at any one time.

It limits the amount of memory which can be locked (lockable memory) to (the available memory indicated at powerup - unlockable\_mem).

If unlockable\_mem is greater than available memory, then none of the memory is lockable.

Lockable memory is used for:

- process images and overhead locked with *plock(2)*
- shared memory segments locked with the SHM\_LOCK command of the *shmctl(2)* system call
- kernel memory allocated for LAN by the *npowerup* command
- miscellaneous dynamic kernel data structures used by the shared memory system and some drivers.

Any call which needs to use lockable memory may fail if its value is too small. Note that lockable memory is a limit on the amount of memory which can be locked, but that this memory is available for virtual memory except when it is locked.

**Cost**

None.

**Dependencies (interactions with other system values)**

`unlockable_mem <= physical memory`

---

## using\_array\_size

### Name

`using_array_size`—size of the diskless cnodes' using array

### Range

0 → memory limited

### Default

0	If the diskless code is not configured into the kernel.
<code>nproc</code>	If the diskless code is configured into the kernel. Unless you change other configurable parameters, this value will be 84.

The diskless code is configured into the kernel if the kernel description file you used contains the driver called `dskless`. Refer to the section on kernel configuration for more information.

### Use

`using_array_size` defines the size of the kernel's using array. The using array is an array of kernel structures that holds information related to outbound requests. It is the same size as the `nproc` configurable parameter.

Each outbound request consumes a single `using_array[]` entry that is held until the request is satisfied. Outbound requests are not discarded if all `using_array[]` slots are utilized; they are put to sleep until the required resource is available.

If both of the configurable parameters, `dskless_node` and `server_node` are equal to zero, the system is treated as a standalone, non-clustered kernel. In this situation, `using_array_size` is zero and `using_array[]` is not compiled into the kernel.



**using\_array\_size**

## **Cost**

16 bytes × using\_array\_size

## **Dependencies (interactions with other system values)**

The size of the using array is directly dependent upon the nproc configurable parameter.

---

## Swap Space Parameter Interaction

If you change `maxdsiz`, `maxssiz`, `maxtsiz`, or `shmmax`, there are four other operating system parameters that you must also change. The four operating system parameters (`dmmin`, `dmmax`, `dmttext`, and `dmshm`) do not have a separate description page in this appendix, but Table D-1 explains the values these operating system parameters must have for given values of `maxdsiz`, `maxssiz`, `maxtsiz`, or `shmmax`. All of these swap space system parameters interact, and a wrong value might make your virtual memory system unworkable.

---

**Note**            The system parameters `dmmin`, `dmmax`, `dmttext`, and `dmshm` should only be changed when it is necessary to change `maxdsiz`, `maxssiz`, `maxtsiz`, or `shmmax`. At all other times, these parameters should be left alone.

---

In the following example, assume you have a stack size equal to 90 Mbytes. Using the information in Table D-1, you would set `dmmin` to 16, `dmmax` to 2048, `dmttext` to 2048, and `dmshm` to 2048. Based on this information, you would need the following entries in your configuration description file (`/etc/conf/dfile`):

```
* tunable system parameters

maxdsiz      116*1024*1024
dmmin        16
dmmax        2048
dmttext      2048
dmshm        2048
```

Refer to the explanation of `dfile` in the “Using the config Command” section of Chapter 6.

## Swap Space Parameter Interaction

Table D-1 shows the values that should be set in order to make the swap space parameters function correctly. The sizes shown for `maxdsiz`, `maxssiz`, `maxtsiz`, and `shmmax` are in Mbytes. The sizes shown for `dmmin`, `dmmax`, `dmttext`, and `dmshm` are in 4 Kbyte disk blocks.

Note that if you have a Model 310 or less than 20 Mbytes of swap space, you may want to reconfigure your kernel with the numbers shown with a superscript 1 (<sup>1</sup>) for optimum use of your swap space.

All numbers shown with superscript 2 (<sup>2</sup>) are the default settings.

**Table D-1. Virtual Memory Parameter Table**

If <code>maxdsiz</code> or <code>maxssiz</code> is:	If <code>maxtsiz</code> is:	If <code>shmmax</code> is:	<code>dmmin</code> should be:	<code>dmmax</code> should be:	<code>dmttext</code> should be:	<code>dmshm</code> should be:
<= 15 MB	<= 10 MB	<= 3 MB	16 <sup>1</sup>	256 <sup>1</sup>	256 <sup>1</sup>	256 <sup>1</sup>
<= 30 MB	<= 20 MB	<= 6 MB	16 <sup>2</sup>	512 <sup>2</sup>	512 <sup>2</sup>	512 <sup>2</sup>
<= 59 MB	<= 40 MB	<= 12 MB	16	1024	1024	1024
<= 116 MB	<= 80 MB	<= 24 MB	16	2048	2048	2048
<= 228 MB	<= 160 MB	<= 48 MB	16	4096	4096	4096
<= 448 MB	<= 320 MB	<= 96 MB	16	8192	8192	8192
<= 880 MB	<= 640 MB	<= 192 MB	16	16384	16384	16384
<= 1728 MB	<= 1280 MB	<= 384 MB	16	32768	32768	32768
<= 3392 MB	<= 2560 MB	<= 768 MB	16	65536	65536	65536
<= 6656 MB	<= 5120 MB	<= 1536 MB	16	131072	131072	131072



# Glossary

---

The following terms and definitions are defined as they apply to the Series 300 HP-UX operating system.

## **address**

In the context of peripheral devices, a set of values which specify the location of an I/O device to the computer. The address is composed of up to four elements: select code, bus address, unit number and volume number. You can read about addresses in the *Peripheral Installation Guide* and in “Adding/Moving Peripherals” in Chapter 5 of this manual.

## **access**

An interaction between a subject and an object that results in the flow of information from one to the other.

## **access control mechanism (ACM)**

An algorithm and data structure that supports access control decisions. It mediates decisions about whether specific subjects can access specific system objects in specific ways. An ACM might be implicit (contextual), e.g. “only user ID 0 can access files in /etc”. It might be explicit, as in a list of objects and their access rights. An explicit ACM might be distributed (data stored with protected objects), as in the UNIX file permissions scheme, or it might be centralized, as in a file describing access rights for all protected system objects.

## **ARPA hostname**

Needed for ARPA Services. A name consisting of one field containing any printable character except spaces, newlines, or the pound sign (#). Assigned in */etc/hosts*.

**attribute**

In an HP-UX cluster, the attributes are characteristics of the cnodes, such as the cnode name or processor type. The attributes are collectively referred to as the cnode's context.

**audit trail**

A set of records that collectively provide documentary evidence of processing, used to aid in tracing from original transactions forward to related records and reports, and/or backwards from records and reports to their component source transactions.

**block**

The fundamental unit of information HP-UX uses for access and storage allocation on a mass storage medium. Block size on an HFS file system can be either 4 Kbytes or 8 Kbytes, and is set at file system creation. The default block size is 8 Kbytes. However, to present a more uniform interface to the user, most system calls and utilities use "block" to mean 512 bytes, independent of the actual block size of the medium.

**block mode**

Buffered I/O. With buffered I/O data is held in the buffer cache, then transferred one block at a time. Block size for buffered I/O is not the same as block size on the file system. Block size for buffered I/O is defined as `BLKDEV_IOSIZE` in `/usr/include/sys/param.h`. Compare with **raw mode**.

**boot or boot-up**

The process of loading, initializing and running an operating system.

**boot area**

The first 8 Kbytes of the disk that are reserved for system use. This area contains the LIF volume header, the directory that defines the contents of the volume, and the bootstrapping program.

**boot ROM**

A program residing in ROM (Read Only Memory) that executes each time the computer is powered-up. The function of the boot ROM is to run tests on the computer's hardware, find some devices accessible through the computer and then load either a specified operating system or the first operating system found according to a specific search algorithm. The bootstrap program uses the boot ROM's mass storage or LAN drivers to

load and pass control to the kernel. When the kernel gains control, it completes the job of bringing up the HP-UX operating system.

Depending on your boot ROM version, the boot ROM displays may differ slightly from those shown in this manual; any differences between boot ROM versions are noted in this manual when the topic in question is discussed. The boot ROM identifies its version when power is applied to the computer.

**bus address**

Part of an address used for devices, especially devices on an HP-IB (HP Interface Bus); a number determined by the switch setting on a peripheral which allows the computer to distinguish between two devices connected to the same interface. A bus address is sometimes called a “device address”, and no two devices on the same HP-IB can have the same bus address.

**bytes per inode**

This specifies the number of data bytes (amount of user file space) per inode slot. The number of inodes is calculated as a function of the file system size. The default value is 2048.

**certification**

The technical evaluation of a system’s security features, made as part of and in support of the approval/accreditation process, that establishes the extent to which a particular computer system’s design and implementation meet a set of specified security requirements.

**CDF**

Context Dependent File. The mechanism used by HP-UX clusters to share a path name between different cnodes in a cluster. A CDF consists of a hidden directory and one or more subfiles.

**client**

A process that is requesting some service from another process. In this manual, client is sometimes used as another name for a diskless cnode.

**cluster**

An HP-UX cluster is one or more workstations linked together with a local area network (LAN), but consisting of only one file system.

**cluster node (cnode)**

A computer in a cluster.

**cluster server process (CSP)**

A special kernel process that runs on a machine in a cluster to satisfy requests from other nodes in the cluster. There are two kinds of CSP, the limited CSP (LCSP) and the general CSP (GCSP).

**cnode**

Cluster *node*. A computer in a cluster.

**connect session**

This denotes the period of time in which a user is connected to the system. It starts when the user logs in and finishes when the user logs out.

**context**

An ASCII string made up of a number of attributes that ordinarily determines which subfile (if any) is accessed from a CDF. Each workstation on a cluster has an associated context which is set at boot time. The context attributes are: cnode name; processor type; floating point hardware; file system location; and the string "default".

**cron**

This process wakes up every minute to execute commands at specified dates and times, according to instructions in files contained in the directory `/usr/spool/cron/crontabs`. See the *cron(1M)* and *crontab(1)* entries in the *HP-UX Reference* for more details.

**CS/80**

A family of mass storage devices that communicate with a computer via the CS/80 (Command Set '80) or SS/80 (Sub Set '80) command set. In HP-UX your file system can be on either a SCSI drive or a CS/80 drive.

**cylinder**

One or more vertical collections of tracks in a disk pack (refer to Figure Glossary-1 later in this section). Disk accesses within a cylinder do not need a seek.

**cylinder group**

One or more consecutive cylinders (refer to Figure Glossary-1 later in this section). Each cylinder group contains a superblock, inodes, cylinder group information, and data blocks.



**cylinder group information**

A data structure located in the cylinder group that contains information about the cylinder group such as numbers of available inodes, data blocks, and fragments, and bitmaps to free space in the cylinder group.

**DASS**

Direct Access Secondary Storage. A new area of the mass storage business structure that falls between online magnetic disk primary storage and offline magnetic tape secondary storage.

**destination device**

The mass storage device on which HP-UX is to be installed or updated. The destination device must be a hard disk drive.

**Discretionary Access Control (DAC)**

A means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that a subject with a certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject (unless restrained by mandatory access control).

**diskless cnode**

A cnode that does not have a local file system. Its file system resides, remotely, on the root server.

**diskless workstation**

The same as a diskless cnode.

**driver number**

Same as major number.

**disk**

The term used for a collection of recording platters contained in a single disk drive. Disk is synonymous with disk pack.

**domain**

The set of objects that a subject has the ability to access. A set of (object, rights) pairs. Each pair specifies an object and some subset of the operations that can be performed on it. A right in this context means permission to perform one of the operations.

## **etc/rc**

This is the system initialization shell script. The actions that it performs depend on the state in which it is invoked. To automatically start System Accounting whenever the system is switched to multi-user mode, a command must be added to *rc*. See the chapter “System Boot and Login” in this manual, and *rc(1M)* in the *HP-UX Reference* for more details on the use of *rc*.

## **etc/shutdown**

A shell script that has the primary function of terminating all currently running processes in an orderly and cautious manner. See *shutdown(1M)* for details on this shell.

## **exchange**

An exchange is when an optical autochanger replaces one disk surface in the optical drive with another.

## **file**

A discrete collection of information described by an inode and residing on a mass storage medium.

## **file types**

Several file types are recognized by HP-UX. The file type is established at the time of the file’s creation. The types are:

- Regular files - Contains a stream of bytes. Characters can be either ASCII or non-ASCII. This is generally the type of file a user considers to be a file: object code, text files, nroff files, etc.
- Directory - HP-UX treats directories like regular files, with the exception that writing directly to directories is not allowed. Directories contain information about other files.
- Block special files - Device files that buffer the I/O. Reads and writes to block devices are done in block mode.
- Character special files - Device files that do not buffer the I/O. Reads and writes to character devices are in raw mode.
- Network special files - contain the address of another system.
- Pipes - A temporary file used with command pipelines. When you use a pipeline, HP-UX creates a temporary buffer to store

information between the two commands. This buffer is a file, and is called a pipe.

- FIFO - A named pipe. A FIFO (First In/First Out) has a directory entry and allows processes to send data back and forth.
- symbolic link - A type of file that indirectly refers a path name.

### **file system**

The organization of files and directories on a hard disk. The HFS file system is an implementation of the HP-UX directory structure.

### **fragment**

A piece of a block. This is the smallest unit of information HFS will read or write. The lower limit of a fragment is DEV\_BSIZE (defined in `/usr/include/sys/param.h`). Fragment size is set at file system creation.

### **free space threshold**

Specifies minimum *percentage* of free disk space allowed. Once the file system capacity reaches this threshold, only the superuser is allowed to allocate disk blocks. The default is 10%; if it is less, file system performance degrades. The free space threshold is set when you create a new file system.

### **HFS file system**

High performance File System. This is the file system implemented on your Series 300 HP-UX system. Other models of Hewlett-Packard computers running HP-UX may have a different file system.

### **hidden directory**

A directory used to implement a CDF. It is called hidden because it is normally treated and seen as a file. It can be accessed as a directory only by appending the special character “+” to its name.

### **HP-UX directory structure**

The hierarchical grouping of directories and files on HP-UX.

### **HP-UX system hostname**

Name you assign to your system (from a line on `/etc/rc`). It is used for UUCP, mail, and other programs.

**inode**

A data structure containing information about a file such as file type, pointers to data, owner, group, and protection information.

**internet address**

An address used by NS-ARPA Services. It consists of two parts: a network number and a host number. Nodes on the same LAN will have the same network number and distinct host numbers.

**ITE**

The Internal Terminal Emulator program which allows a bit-mapped display to function as a standard computer terminal.

**kernel**

The core of the HP-UX operating system. The kernel is the compiled code responsible for managing the computer's resources; it performs such functions as allocating memory and scheduling programs for execution. The kernel resides in RAM (Random Access Memory) whenever HP-UX is running.

**LIF**

Logical Interchange Format. LIF is Hewlett-Packard's standard file format, used for transferring files between Hewlett-Packard systems. Since LIF is a standard, files with LIF format can easily be transferred between different Hewlett-Packard computers (refer to the *LIF(4)* entries in the *HP-UX Reference*).

**link level address**

A unique 12-digit hexadecimal number which is part of every LAN card. This number appears on the LAN card hardware, on the boot ROM screen, and can be obtained using the `landiag` program.

**login**

The process of a user gaining access to HP-UX. This process consists of entering a valid user name and its associated password (if one exists).

**magneto-optical**

Magneto-optical is a form of rewritable optical technology.

**major number**

An index into a table in the kernel. It is needed to use peripheral devices. The major number for a particular device can be found in the tables in “Adding/Moving Peripheral Devices” in the “Customizing” chapter.

**mechanical picker or mechanical changer**

The part of the optical autochanger that moves optical disks from slot to drive and vice versa.

**media**

In terms of optical products, this is the optical disk that holds the data. The term includes the plastic cartridge that houses the optical disk.

**minor number**

A hexadecimal number made up of a select code and other information specific to the peripheral device you are setting up. Refer to “Adding/Moving Peripheral Devices” in the “Customizing” chapter for information on minor numbers.

**MO**

See **magneto-optical**

**multi-user run-level**

A run-level of HP-UX when terminals in addition to the system console allow communication between the system and its users. The multi-user run-level (not to be confused with a multi-user system) is run-level 2 as shipped.

**MUX**

MUX is an abbreviation for Asynchronous Multiplexer. The HP 98642 four channel MUX is available for the Series 300. Each channel is an RS-232C port which is normally associated with a /dev/ttyXX file.

**named object**

Objects which have names, are visible at the TCB interface, and are shared among users.

**NCSC**

National Computer Security Center, the government agency that wrote the guidelines for trusted systems.

**node**

A computer on a network.

**NS\_ARPA Services**

A combined networking product providing both NS and ARPA services. The NS\_ARPA Services networking product enables your Series 300 to transfer files to and from remote hosts, log into remote hosts, execute commands on remote hosts, and send mail to and receive mail from remote hosts on the network.

**NS nodename**

Needed for Network Services and the rlb daemon. A name consisting of three fields separated by periods, i.e. node.domain.organization. Each field can contain up to 16 alphanumeric, case insensitive characters.

**objects**

A passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are: records, blocks, pages, segments, files, directories, directory trees, and programs, as well as bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, networks nodes, etc.

**optical autochanger**

A rewritable optical mass storage peripheral which includes the mechanics to move optical disks in and out of drive(s), the drive(s), media and controller electronics.

**password**

A private character string that is used to authenticate an identity.

**path name**

A series of directory names separated by / characters, and ending in a directory name or a file name.

**primary storage**

Typically consists of fixed hard disk(s), used for fast, random-access applications. The primary storage devices are used as online system disks.

**process**

A process is the environment in which a program (or command) executes. It includes the program's code, data, status of open files, and value of variables. It is completely characterized by a single current execution point and address space. For example, whenever you execute an HP-UX command, you are creating a process; whenever you log in, you create a process. For additional information on processes, read the chapter "Concepts."

**raw mode**

Unbuffered I/O. Data is transferred directly between the device and the user program requesting the I/O, rather than going through the file system buffer cache. Compare with **block mode**.

**rewritable optical**

An optical disk technology which can be repeatedly written.

**resource**

Anything used or consumed while performing a function. The categories of resources are: time, information, objects (information containers), or processors (the ability to use information). Specific examples are CPU time; terminal connect time; amount of directly-addressable memory; disk space; number of I/O requests per minute.

**root**

Root refers to the highest level directory (root directory or /).

**root server**

The node in an HP-UX cluster with the local file system, capable of supporting other diskless workstations.

**secondary storage**

The storage device(s), typically tape drives, used to back up and archive data stored on the system disks (primary storage). Secondary storage is also used to log transaction, interchange data and distribute software. Secondary storage devices always use removable media.

**security policy**

The set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information.

**select code**

Part of an address used for devices; a number determined by switch settings on the interface card. The select code determines the interface card's location in the processor address space. Each interface card is in turn connected to a peripheral. Multiple peripherals connected to the same interface card share the same select code.

**sensitive information**

Information that, as determined by a competent authority, must be protected because its unauthorized disclosure, alteration, loss, or destruction will at least cause perceivable damage to someone or something.

**shell**

A program that interfaces between the user and the operating system. HP-supported shells are:

- /bin/sh
- /bin/csh
- /bin/ksh
- /bin/rsh
- /bin/rksh
- /bin/pam

**source device**

The mass storage device from which HP-UX is installed. The source device must be a cartridge tape drive or flexible disk drive.

**special file**

Often called a device file, this is a file associated with an I/O device. Special files are read and written just like ordinary files, but requests to read or write result in activation of the associated device. These files normally reside in the /dev directory.

**standalone**

A machine which is not part of an HP-UX cluster.

**subfile**

Part of a CDF in an HP-UX cluster. The subfiles are under the hidden directory, and are named for one of the system's context attributes.



**subject**

An active entity, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. Technically a process/domain pair.

**superblock**

A data structure containing global information about the file system such as file system size, disk information, and cylinder group parameters. The superblock is created at the same time as the file system and is replicated into each cylinder group.

**surface**

In terms of optical disks, this is one of the disk sides—surface 1 or 2.

**system administrator run-level**

A run-level of HP-UX when the system console provides the only communication mechanism between the system and its users. Init state s is the system administrator run-level.

**system console**


A keyboard and display (or terminal) given a unique status by HP-UX and associated with the special device file `/dev/console`. All boot ROM error messages (messages sent prior to loading HP-UX), HP-UX system error messages, and certain system status messages are sent to the system console. Under certain conditions (for example, the single-user state), the system console provides the only mechanism for communicating with HP-UX.

**TCSEC**

Trusted Computer Systems Evaluation Criteria, also known as the “Orange Book”. This is the book where evaluation criteria for trusted systems is documented.

**track**

One of several concentric circles on the surface of a disk upon which data is recorded (refer to Figure Glossary-1).



**Figure Glossary-1. Track, Cylinder, and Cylinder Group on a Disk**

**trap door**

A hidden software or hardware mechanism that permits system protection mechanisms to be circumvented. It is activated in some non-apparent manner (e.g., special “random” key sequence at a terminal).

**Trojan horse**

A computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the

legitimate authorizations of the invoking process to the detriment of security. For example, making a “blind copy” of a sensitive file for the creator of the Trojan Horse.

**trusted computer system**

A system that employs sufficient hardware and software integrity measures to allow its use for processing simultaneously a range of sensitive or classified information.

**trusted computing base (TCB)**

The totality of protection mechanisms within a computer system—including hardware, firmware, and software—the combination of which is responsible for enforcing a security policy. A TCB consists of one or more components that together enforce a unified security policy over a product or system. The ability of a TCB to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user’s clearance) related to the security policy.

**trusted process**

A process that is restricted and is only run by a user with appropriate privilege. For HP-UX 6.5 a trusted process is a process with effective user ID of 0 (superuser).

**trusted software**

The software portion of a Trusted Computing Base.

**unit number**

Part of an address used for devices; a number whose meaning is software- and device-dependent but which is often used to specify a particular disk drive in a device with a multi-drive controller. When referring to single-controller integrated disk/tape or disk/flexible disk drive, a unit is used to distinguish between disk and cartridge tape drives or hard disk and flexible disk drive.

The unit number also selects a single partition on the 913x series.

**user**

Any person who interacts directly with a computer system.

**UUCP**

Unix-to-Unix File Copy Package, mainly used for mail transport

**verification**

In a trusted system, the process of comparing two levels of system specification for proper correspondence (e.g., security policy model with top-level specification, TLS with source code, or source code with object code). This process may or may not be automated.

**volume number**

Part of an address used for devices; a number whose meaning is software- and device-dependent but which is often used to specify a particular volume on a multi-volume disk drive. The volume number is also used to inform the device driver of special handling semantics (such as printer drivers skipping over perforations).

# Index

---

## A

- access control lists 2-33, 8-10
- access permission 2-32, 2-35
- ACL 2-33
- adding a diskless cnode 4-92
- adding users 5-3, 5-12
- adding users to a group 5-13
- addressing information 4-17
- administering an HP-UX cluster 2-74
- advisory locks 2-38
- alternate superblock A-4, A-8, 2-16, 2-26
- amigo 4-45
- applications 6-17
- archive backup 5-21
- archiving programs 5-71
- ARPA hostname 4-86, 4-96
- assisting other users 1-9
- attended mode 3-11, 3-13, 4-134, 6-6
- audit file overflow 8-17, 8-28
- auditing 8-11
- auditing tasks 8-20
- autoboot 3-27
- autochanger cartridge tape drive 5-30
- autocreation 2-84
- automatic mounting of all file systems 4-146
- automatic unmounting of all file systems 4-146
- available memory 2-63

## B

- backing up the system
  - archive backup 5-21
  - automatic backups 5-48

- backing up selected files 5-41
- backup device file 5-30
- description 5-21
- error message 5-31
- file restoration 5-43
- HP-UX cluster 5-21, 5-26
- incremental backup 5-21
- media 5-26
- problems B-10
- scripts 5-27
- system clock considerations 5-22
- times 5-27, 5-29
- using `cron` 5-48
- backup kernel 2-15, 3-14
- backup kernel in an HP-UX cluster 6-6, 6-37, 6-41, 6-51
- backup script 5-27
- back-up system administrator 1-9, 2-12
- backupf script 5-27
- bad logins 3-24
- baud rate 4-34
- bcheckrc A-1, A-3
- BFS file system 2-31
- bit-mapped display 3-10
- bits per inch (bpi) 2-43, 2-46
- block device file 2-20, 2-68, 4-13, 4-16
- block major number 4-17
- boot area 2-15
- boot ROM
  - attended mode 3-11, 3-13, 4-134, 6-6
  - autoboot 3-27
  - booting across LAN 3-27
  - general description 3-2
  - input device 3-8
  - loading HP-UX 3-7, 3-9, 3-11
  - output device 3-8
  - search sequence 3-13
  - system console 3-8
  - unattended mode 3-11-12, 3-14
- booting the recovery system 4-130, 4-133
- bootup
  - automatic mounting of file systems 4-146

- booting a diskless cnode 3-6, 3-15
- booting a new operating system 3-30
- booting an HP-UX cluster 3-5
- booting the system 3-1-2, 3-9, 3-12
- error messages B-5
- `fsck` A-1, A-26, 3-4, 3-17, 3-27
- loading the operating system 3-12
- operating system selection menu 3-14
- problems B-6, 3-27, 3-32, 4-126, 4-129, 5-45

Bourne shell 5-8

buffer cache 2-23, 2-40

## C

- C shell 5-8
- calculating swap space 6-17
- cartridge tape, description 2-42
- cartridge tape drive
  - autochanger 5-30
  - for backups 5-26
  - setting up 4-61
  - updating or installing software 4-196

CDF

- context attributes 2-84
- creating 2-82
- description 2-78
- examining 2-83
- examples 2-85
- finding 2-82
- implementation 2-80
- mixing context attributes 2-83
- removing 2-82
- structure 2-80
- system CDFs 2-101
- using 2-79

- changing a password 5-50
- changing run-levels 3-19, 3-29, 5-51
- changing the system clock 5-68
- character device file 2-20, 2-68, 4-13, 4-16
- character major number 4-17
- characters per inch (cpi) 2-46
- charging fees 7-12, 7-50

- checking the file system (see `fsck`) 5-1
- `checklist`, programs that use 4-76
- clean byte A-2, 3-17, 4-140
- cluster server process ( see HP-UX cluster - CSP) 4-86
- clustered environment 4-88
- cnode, adding a diskless cnode 4-92
- cnode hostname 4-86
- cnode name 4-81, 4-96
- code segment 2-59-60, 2-64
- commands
  - `config` 2-60, 2-64
  - `fsck` 2-16, 2-24, 2-26
  - `fsdb` 2-16, 2-31
  - `tset` 4-73
- communicating with users 5-56
- communication protocol 3-23
- components of system security 8-3
- `config` 2-60, 2-64, 4-16, 6-2
- configuring
  - configuration description file 6-3, 6-5, 6-8-9, 6-11, 6-13, 6-22, 6-25-26, 6-29
  - description 4-149
  - device drivers 6-8, 6-33
  - HP-UX cluster kernels 6-52
  - interface cards 6-8
  - kernel drivers 4-16
  - list of drivers 6-12
  - operating system parameters D-1, 6-4, 6-29
  - `reconfig` 4-149, 6-33
  - swap space 6-4, 6-13
- connect session accounting 7-10, 7-23
- connecting peripherals 4-15
- context 2-77
- context attributes 2-77
- continuation inode A-20, A-25, 2-27
- continuous process 3-24, 5-51
- conventions 1-3
- copying HP-UX files to LIF volumes 5-73
- correcting file system corruption A-1, 2-25
- `cpio` B-10, 5-70
- crash A-1, 3-32, 5-21
- creating a login directory 5-13



- creating a new file system 4-107
- creating a password 5-16
- creating a recovery system 4-129
- creating LIF volumes 5-64
- creating swap space 6-15
- creating system run-levels 4-3, 4-125
- cron 5-48, 5-67, 7-4
- cron in an HP-UX cluster 2-105
- CS/80 4-45
- .cshrc 4-7
- CSP 2-96
- customization
  - config vs. reconfig 4-149, 6-2
  - description 4-1
  - environment files 3-20
  - HP-UX cluster 2-96
  - kernel 6-1
  - reconfig 4-12, 4-149
- cylinder group 2-14, 2-16, 2-20, 2-22
- cylinder group information 2-16, 2-24

## **D**

- data blocks 2-16, 2-24, 2-30
- data communication 4-69
- data loss 5-21
- data safe 5-21
- data segment 2-59-60, 2-64
- data size D-96
- data storage 2-20
- date 3-5, 4-10, 5-66
- dd command 5-73
- deadlock 2-40
- default files 4-13
- demand loading 2-59, 2-67
- demand paged virtual memory management 2-58, 2-63
- dependent fileset C-2, C-9
- detecting file system corruption A-1, 2-25
- determining amount of swap space 6-16
- device classes 2-68, 4-12
- device drivers
  - configuring 6-2, 6-8

- definition 2-68
- device file
  - block 2-20, 2-68, 4-74
  - block vs. character 4-13, 4-16-18, 4-107
  - character 2-20, 2-68, 4-30, 4-36, 4-63, 4-67
  - creating 4-15
  - default 4-13
  - definition 2-68
  - description 4-12
  - digitizer 4-66
  - file system mounting 4-141
  - GPIO 4-68
  - HP-HIL 4-68
  - mass storage 4-39, 4-56, 4-61
  - miscellaneous device class 4-18
  - modems 4-30
  - naming convention 4-17, 4-30, 4-36, 4-39, 4-56, 4-61, 4-63, 4-67
  - network 2-68
  - path name 4-17
  - plotter 4-66
  - pseudo devices 4-36
  - required for the system 4-18, 4-44
  - terminals 4-30
  - type 4-36
- device I/O 2-68
- df command 5-58
- digitizer 4-66
- Direct Access Secondary Storage 2-52
- direct connect port 4-32
- directories
  - /lost+found A-22
  - moving 5-77
  - protection 2-34
  - structure 2-2
- disk
  - crash 3-32
  - creating file system on 4-107
  - /etc/disktab file 4-107
  - initializing 4-107
  - interface 4-45
  - layout 2-14

- removing from the system 4-77
- usage 5-1, 5-58
- disk drive
  - HP-UX cluster 4-43
  - setting up 4-39
- disk space allocation 2-21
- disk space usage accounting 7-10, 7-17
- diskless cnode, adding 4-92
- drivers
  - bundled 4-14
  - configuring 4-16
  - default system 4-14
  - /etc/master file 4-14
  - list of 6-12

## E

- effective group ID 2-10
- effective user ID 2-10, 2-12
- enforcement locking mode 2-35, 2-39
- ENOMEM 2-60, 2-64
- .environ 4-7
- environment file
  - .cshrc 4-7
  - .environ 4-7
  - /etc/btmp 4-6
  - /etc/checklist A-2, A-6, 4-8, 4-141, 4-145
  - /etc/csh.login 3-25, 4-10
  - /etc/group 4-4
  - /etc/inittab A-1, 3-16, 3-23, 4-3, 4-32, 4-125
  - /etc/issue 4-10
  - /etc/motd 4-4
  - /etc/passwd 4-4
  - /etc/profile 3-25, 4-5, 4-10
  - /etc/rc 3-20, 4-4, 4-10
  - /etc/securetty 4-6
  - /etc/ttytype 4-11, 4-32
  - /etc/utmp 4-6
  - /etc/wtmp 4-5
  - .exrc 4-8
  - general information 4-3
  - local 4-7

.login 4-7  
.profile 4-7  
/usr/lib/terminfo 4-8  
environment variable  
HOME 4-7, 5-6  
PATH 3-25, 4-5, 4-7  
PS1 4-7  
PS2 4-7  
TERM 3-25, 4-29  
TZ 3-25, 4-5, 4-10, 5-66  
environment variables  
TERM 4-71  
error message  
<command>: not found B-11  
cpio: Out of phase B-10  
danger: mfree map overflow B-3  
dealing with the file system B-4  
during backup 5-31  
during bootup B-5  
general B-1  
<name>:table is full B-3  
no more processes B-3  
No Shell B-6  
Not enough memory B-5  
panic B-2  
pid <num>:xalloc B-3  
Sorry, pid *pidnum* was *message* 6-13  
swap: rmap ovflo, lost<blocknumber,blocknumber> 6-15  
system panic 3-32  
Unable to access ttyxx B-13, 4-29, 4-35  
WARNING: avail mem too low, may deadlock B-5  
/etc/bcheckrc 3-16  
/etc/btmp 4-6  
/etc/checklist A-2, A-6, 4-8, 4-74, 4-141, 4-145  
/etc/clusterconf 2-99, 4-85  
/etc/csh.login 4-10  
/etc/disktab 4-107-108, 6-27  
/etc/fsck A-1  
/etc/group 4-4  
/etc/init 3-7, 3-16  
/etc/inittab A-1, 3-7, 3-16, 3-23, 4-3, 4-32, 4-126

- /etc/issue 4-10
- /etc/master 4-14
- /etc/motd 4-4
- /etc/passwd 4-4
- /etc/profile 4-5, 4-10
- /etc/rc 3-20, 4-4, 4-10, 7-3, 7-30
- /etc/rcflag 3-20
- /etc/sbtab A-4, A-8
- /etc/securetty 4-6
- /etc/ttytype 4-11, 4-32
- /etc/update C-15
- /etc/utmp 4-6
- /etc/wtmp 4-5
- evaluating users' needs 1-6, 1-10
- execute permission 2-32, 2-35
- .exrc 4-8

## F

- FIFO 2-19

## FIGURES

- Add a New User Menu 5-6
- Adding/Removing Cnode Main Menu 4-95, 4-104
- Boot ROM Display 3-3
- Boot-Up Sequence 3-8
- CDF containing multiple context attributes 2-89
- CDF Structure Using Different Floating Point Hardware 2-86
- CDF with localroot and remoteroot Context Attributes 2-88
- CDF with Mixed Attributes 2-88
- CDF with Mixed Context Attributes 2-83
- CDF with the "default" Attribute 2-89
- Change Destination Directory Menu 4-200, 4-217
- Change Source Device Menu 4-202, 4-218
- Choosing a Group Name 5-7
- Choosing a Shell 5-8
- Choosing Interface Card Kernel Device Drivers 6-48
- Choosing Kernel Device 6-46
- Choosing Line Printer Interfaces 4-162
- Choosing Line Printer Model 4-161
- Choosing Optional Software Products 6-50
- Choosing Port Usage 4-26
- Choosing Terminal Speed 4-25

- Configuring Terminal Support Menu 4-24
- Creating a Cluster Menu 4-85
- Custom Operating System Menu 6-44
- Delete a User Menu 5-18
- Devices in Custom Operating System Menu 6-45
- /etc/inittab in the example cluster 2-79
- /etc/rc functions 3-21
- Example Boot ROM Display 4-93
- Example Cluster 2-73
- Example /etc/checklist File 4-77
- File Permission Bits 2-33
- File Permission Bits of 6754 2-36
- File Permissions Bits of 0754 2-34
- File System Hierarchy After Unmounting 4-144
- File System Hierarchy Before Mounting 4-142
- File System Hierarchy While File System is Mounted 4-143
- File System Layout on Disk 2-14
- Fileset Menu for Update on Flexible Disk 4-222
- Fileset Menu for Updates on Cartridge Tape 4-208
- Flow of /etc/bcheckrc Program 3-18
- Flow of /etc/brc Program 3-19
- Flow of /etc/rc Program 3-22
- Form for Adding a Cluster Node 4-96
- Form for Deleting a Cluster Node 4-105
- HP Series 6300 Model 20GB/A - Optical Autochanger 2-55
- HP-UX Partitions, Filesets and Files C-2
- HP-UX Tree Structure for Major Directories 2-2
- Inode Addressing Example 2-19
- Interface Cards Menu 6-47
- Line Printer Configuration Menu 4-160
- Line Printer Spooling 4-154
- LP Spooler Directories and Files 4-177
- Magnetic Tape Format 2-44
- Main Utility Menu 4-199, 4-215
- Messages during Bootup 3-4
- Nested CDFs 2-87
- Optional Software Products Menu 6-49
- Partition Menu for Update on Cartridge Tape 4-206
- Reading File Menu 4-203, 4-219
- Reconfig and Config Functionality Overlap 6-2
- Reconfig Main Menu 4-22, 4-84, 4-94, 4-103, 4-151, 4-158, 5-4, 6-35, 6-39, 6-42

- Reconfig Tasks 4-149
- Reconfig User Configuration Menu 5-5
- Reconfig Workstation Configuration Menu 4-23, 4-159, 6-40
- Regular File Mapping Scheme and the Inode Structure 2-18
- Structure of a CDF 2-81
- Swap Space 6-25
- System Accounting Data Flow Diagram 7-14
- System Accounting Directory Structure 7-16
- Terminal and Modem Connections 4-20
- The Four Customization Programs 4-2
- The Kernel as a CDF 6-52
- User Process Logical Address Space 2-61
- Workstation Configuration Menu 6-36, 6-43
- Write Protect Mechanism on Cartridge Tape 4-197
- Write Protect Mechanism on Flexible Disk 4-212
- file owner 2-32, 2-35
- file restoration 5-43
- file system
  - access mode A-7
  - alternate superblock A-4
  - alternate superblock location 4-109
  - alternate superblock locations 2-16
  - automatic mounting/unmounting 4-146
  - backup 5-21
  - boot area 2-15
  - buffer cache 2-23, 2-40
  - checking A-1-2, 3-4, 3-17, 4-140, 5-1
  - corruption 2-25, 2-31
  - crash 3-32
  - creating 4-107, 6-27
  - data storage 2-20
  - debugger (`fsdb`) 2-16, 2-31
  - default 6-14
  - directory structure 2-2
  - disk layout 2-14
  - disk usage 5-1, 5-58
  - errors A-1, B-4
  - `/etc/disktab` file 4-108
  - free block map A-23
  - free space 4-109
  - hierarchy 2-2, 3-16

- logging files 5-59
- memory management 2-58
- minfree threshold 5-59
- mkfs 4-108
- mounting 4-76, 4-109, 4-140–141, 4-145
- mounting/unmounting in an HP-UX cluster 2-97
- newfs 4-108
- restoring from a backup 5-43
- root 3-16–17, 4-8, 4-44, 4-74, 4-77
- size 2-26, 5-58
- swap space 4-74, 4-109
- unmounting 4-140, 4-144
- using the recovery system to repair 4-129

files

- format and compatibility 2-31
- how to list files contained on your installation media C-17
- locking 2-35, 2-38, 2-40
- logging 5-59
- protection 2-32, 3-24
- relation in fileset C-1
- required for the system 4-18
- sharing 2-38
- transferring 5-70

fileset

- choosing which filesets to install C-3, C-9
- definition C-1
- dependent C-2, C-9
- mandatory C-2
- removing 4-148
- size C-9
- tasks you can perform with a fileset C-9

flexible disk drive

- for backups 5-26
- hard error B-11
- setting up 4-15, 4-56
- updating or installing software 4-212

formatting media 4-107, 5-61

fragment 2-16, 2-18, 2-21

frame 2-46

free block A-23, 2-27

free space 2-20



## **fsck**

- alternate superblock locations A-4, 2-16, 2-26, 4-109
- description A-1, A-5, 2-26, 5-1
- during bootup 3-4, 3-17, 3-27
- errors A-5
- file systems to check 4-76
- mounted file system 2-24
- preen mode 3-17
- rebooting A-26

## **fsclean** A-2

## **G**

- gateway 4-92, 4-99
- getty 3-7, 4-3
- good logins 3-24
- GPIO 4-68
- graphics tablet interface 4-15
- group
  - access list 2-11
  - adding users 5-13
  - changing 5-54
  - creating 5-54
  - ID 2-32, 2-35
- group name 5-7

## **H**

- halting the system 3-29
- hard disk drive
  - setting up 4-39
  - system (root) 4-14
- hardware
  - failure 2-25
- hardware address 4-14
- HFS file system 2-13, 2-23, 2-31
- hidden directory 2-81
- HOME directory 5-6
- hostname 4-81, 4-86
- HP-HIL 4-68
- /hp-ux 2-15
- HP-UX cluster
  - adding a diskless cnode 4-92

- after setting up a clustered environment 4-91
- autocreation 2-84
- backing up the system 5-21, 5-26
- backup kernel 6-6, 6-37, 6-41, 6-51
- boot daemon B-43
- booting problems B-18, B-20, B-23, B-26, B-38
- booting the cluster 3-1, 3-5-6, 3-12, 3-15
- cannot log into the diskless cnode B-55
- CDF 2-78-80, 2-82-83, 2-85, 2-101
- changing run-levels 4-125, 5-51
- checking the file system A-1
- cluster configuration problems B-32
- clustered environment 4-88
- cnode name 4-96
- commands and files used in clusters 2-97
- configuring 4-78
- configuring kernels 6-3, 6-52
- context 2-77-78, 2-84
- creating a device file 4-43-44, 4-60, 4-65
- cron 2-105
- CSP 2-92, 2-96, 4-86
- customizing 2-96
- cwall command 5-56
- /etc/clusterconf file 2-99, 4-85
- finding processes 4-145
- hidden directory 2-81
- hostname 4-81, 4-86
- internet address 4-81, 4-86
- LAN gateway 4-92, 4-99
- LAN problems B-29, B-35, B-52
- link level address 4-82, 4-86
- logging files 5-59
- LP spooler 2-105, 4-155, 4-157, 4-184
- mail 2-105
- mounting/unmounting 2-97
- multiple LAN cards 4-89, 4-99
- network files 4-99
- network ifconfig statements 4-90, 4-100
- network route statements 4-90, 4-100
- NFS 2-96
- NS-ARPA files 4-88

- operating system parameters 6-53
- overview 2-69, 2-71
- panic, retry, and other system messages B-46
- power failure 3-32
- process IDs 2-94
- removing a cnode 4-102
- renaming a cnode 4-102, 4-106
- shutdown 3-28, 3-31
- single-point system administration 2-74
- swapping 2-95, 6-56
- system accounting 2-105, 7-4, 7-17, 7-19, 7-22, 7-77
- system administration 2-74, 2-97
- system CDF 4-4-6, 4-9, 4-13, 4-17, 4-28, 4-30, 4-33-34, 4-39, 4-56, 4-61
- system panic 3-31
- troubleshooting B-17, B-41
- updates 4-188, 4-195
- UUCP 2-105
- where to execute commands 2-75
- why use a cluster 2-69
- write command 5-56
- HP-UX hostname 4-81, 4-86, 4-96
- HP-UX operating system
  - customization 4-149

## I

- IDs 2-10
- ifconfig 4-90, 4-100
- immediate report 2-48
- immediate response mode 2-48
- incremental backup 5-21
- initdefault 4-125
- initializing media 4-107, 5-61
- inode 2-16-18, 2-20, 2-24, 2-27
- installation
  - choosing filesets C-1
  - installing the system 4-129
  - optional software 4-186, 4-196, 4-212
- interface card
  - boot ROM search sequence 3-10
  - configuring 6-2, 6-8
  - for printers 4-162

- for terminals or modems 4-25
- HP-IB 4-14, 4-162
- installing 4-13
- line control switch pack settings 3-10
- root disk 4-44
- SCSI 4-45
- select code 4-14, 4-192, 4-203, 4-219
- tape drive 4-66
- internet address 4-81, 4-86
- interrupt characters 2-9

## **K**

- kernel
  - configuring drivers 4-16, 6-2, 6-8
  - configuring interface cards 6-8
  - configuring operating system parameters 6-29
  - configuring swap space 6-13
  - creating 6-38, 6-41
  - customization 6-1
  - drivers 4-14
  - drivers in a fully loaded file system 6-38
  - drivers in minimal operating system 6-34
  - errors B-2
  - initialization procedures 6-15
  - list of drivers 6-12
  - new kernel won't boot 6-6
  - updating 4-186, 4-191
- kernel driver
  - interface card 4-66

## **L**

- LAN
  - gateway 4-89, 4-92, 4-99
  - link level address 4-82, 4-86
  - multiple cards 4-89
  - multiple LAN cards in an HP-UX cluster 4-99
  - network files 4-99
- LIF C-17, 2-15, 5-61, 5-64, 5-73
- line printer spooler 4-154, 4-156
- link level address 4-82, 4-86
- listing files contained on your installation media C-17

- loading HP-UX 3-1-3, 3-9, 3-11
- local area network (LAN) 4-69, 5-77
- local environment files 3-26, 4-7
- local swapping 2-95, 6-56
- lockable memory 2-63
- logging files 5-59
- logical address space 2-58
- .login 4-7
- login
  - after bootup 3-5
  - creating a directory 5-13
  - creating a password 5-14
  - error messages B-6
  - history 4-5
  - process 3-24
  - root 4-6
  - shell 3-25, 5-8
  - system won't accept a password B-7
- /lost+found A-4, A-19, A-22, 2-28, 2-31
- lp 4-156
- LP spooler 4-157
- LP spooling in an HP-UX cluster 2-105, 4-155

## **M**

- magnetic tape 2-42
- magnetic tape drive 4-63
  - for backups 5-26
- mail 5-56
- mail in an HP-UX cluster 2-105
- major number 4-17
- man pages 4-9
- mandatory partition C-2
- mass storage media
  - device file type 4-107
  - formatting 5-61
  - initializing 5-61
  - setting up 4-39, 4-56, 4-61
- master pseudo device 4-36
- MAXBPG 2-22
- mediainit command 5-62
- memory management

- logical address space 2-58
- memory management unit (MMU) 2-58
- overview 2-58
- physical memory utilization 2-62
- message-of-the-day 5-56
- minor number 4-14, 4-17
- mkfs 4-108
- mknod 4-13, 4-16, 4-156
- mkrfs 4-130, 4-132
- modem 4-19, 4-30
- modifying the kernel with an update 4-191
- mounting a file system 4-76, 4-109, 4-140-141, 4-145
- mounting/unmounting a file system in an HP-UX cluster 2-97
- moving between run-levels 4-127, 5-53
- moving directories 5-77
- moving files 5-70
- multiple LAN cards in an HP-UX cluster 4-89

## **N**

- naming convention 4-192
- network device file 2-68, 4-17
- network `ifconfig` statements 4-90, 4-100
- network `route` statements 4-90, 4-100
- networking in a trusted system 8-2, 8-31
- newfs 4-108, 6-27
- news 5-56
- NFS in an HP-UX cluster 2-96
- nine-track tape
  - HP-UX cluster 4-65
- nine-track tape drive (see magnetic tape) 4-15
- non-HP terminals 4-150, 4-189

## **O**

- on-line documentation 4-9
- open files (see processes) 2-9, 4-145
- operating system
  - creating 6-38, 6-41
  - fully loaded 6-38
  - minimal 6-34
  - updating 4-197
- operating system parameters

- argdevnblks 2-65
- configuring 6-4, 6-29
- dmmx D-96-97, 2-64, 6-15
- dmmin D-96-97, 2-64
- dmshm D-96
- dmtxt D-96
- HP-UX cluster 6-53
- maxdsiz D-96-97, 2-60
- maxssiz D-96-97, 2-60
- maxtsiz D-96-97, 2-60
- nbuf 2-41
- ntext 2-66
- reference pages D-1
- shmmax D-96
- shmmaxaddr 2-60
- shmseg 2-60
- unlockable\_mem 2-63
- user process size limits 6-31
- optional filesets and products 4-148

## **P**

- pageout daemon 2-58, 2-63
- PAM 5-8
- parent process 2-6, 2-8
- partition
  - choosing which filesets to install C-3, C-9
  - definition C-1
  - mandatory C-2
  - removing 4-148
- password 8-34
  - changing 5-50
  - creating 5-16
  - forgotten 5-50
  - root 4-129
  - security 5-14
  - setting 5-16
  - system won't accept during login B-7
- password file in a trusted system 8-32
- PATH 3-25, 4-5, 7-5
- peripherals
  - connecting 4-15

- description 4-12
- phantom record 2-49
- physical memory 2-58, 2-63
- physical security 8-3
- pipes 2-19
- plotter 4-66
- port address for terminals and modems 4-31
- power failure 3-32
- power failure on an HP-UX cluster 3-32
- powering down the computer 3-28
- powering up the system 3-1-2, 3-9, 3-12
- preen mode 3-17
- primary inode 2-27
- primary superblock 2-16, 2-26, 3-17
- prime time 7-7
- printer
  - HP-UX cluster 4-155, 4-157, 4-184
  - LP spooler 4-156
  - LP spooler commands 4-172
  - model 4-161, 4-181
  - printer classes 4-179
  - setting up 4-14, 4-154, 4-157
  - spooling 4-183
- privileged group 2-12
- procedural security 8-4
- process accounting 7-8, 7-11, 7-30, 7-35
- process size D-96
- processes
  - configurable operating system parameters 6-30
  - creation 2-6
  - dmmx D-97
  - dmmin D-97
  - dmshm D-97
  - dmtxt D-97
  - /etc/update C-15
  - init 2-8
  - maxdsiz D-97
  - maxtsiz D-97
  - nproc D-53
  - open files 2-9, 4-145
  - process groups 2-8



- process ID 2-6, 2-8, 2-10
- process IDs in an HP-UX cluster 2-94
- shmmax D-97
- terminal affiliation 2-9
- termination 2-8, 2-10
- user process size limits 6-31
- .profile 4-7
- protection 4-125, 5-3, 5-14
- protocol, communication 3-23
- pseudo devices
  - description 4-36
  - kernel driver 6-8

## **Q**

- quiescent state 2-26

## **R**

- read permission 2-32, 2-35
- real group ID 2-10
- real user ID 2-10, 2-12
- reconfig
  - adding a diskless cnode 4-92, 4-94
  - adding kernel drivers 4-16, 6-2, 6-33
  - adding peripheral devices 4-12
  - configuring an HP-UX cluster 4-78
  - general description 4-149
- recovering lost data 5-21
- recovery from a power failure 3-32
- recovery system
  - booting 3-30
  - creating 4-129
  - device files 4-130
  - shutting down 4-135
  - using manually 4-135
- remote swapping 2-95, 6-56
- removing a cnode 4-102
- removing a terminal 4-35, 5-20
- removing a user 5-3, 5-17, 5-20
- removing optional filesets and products 4-148
- renaming a cnode 4-102, 4-106
- respawn 3-23

- RJE 4-69
- root device 4-14
- root disk 4-44
- root file system A-2, 3-16-17, 4-8, 4-74, 4-77, 4-130
- root user 2-12, 4-6, 4-129, 5-51
- route statements 4-100
- run-level
  - 0 4-125
  - 2 3-23
  - changing 3-19, 3-29, 5-51
  - changing in an HP-UX cluster 4-125, 5-51
  - creating 4-125
  - description 3-16, 4-3
  - initdefault 3-19
  - moving between 5-53
  - s 2-26, 3-26, 3-28-29, 4-125, 5-51

## **S**

- sbtabs 4-130
- scsi 4-45
- SDF file system 2-31
- search sequence 3-13
- secondary loader 2-15
- security 8-1, 8-33
  - access control lists 8-10
  - ACL 8-10
  - auditing 8-11
  - components 8-3
  - converting to a trusted system 8-7
  - description 5-3, 5-14
  - logins 3-24
  - physical 8-3
  - policy 8-5
  - procedural 8-4
  - self-auditing 8-17
  - system 8-4
  - TCB 8-9
  - trusted computing base 8-9
- Security note 4-129, 4-196, 5-3, 5-25
- select code 4-14
- self-auditing 8-17

- serial interface card 3-10
- set group ID bit (`sgid`) 2-35, 2-39
- set user ID bit (`suid`) 2-35
- setting the date 4-10, 5-67
- setting the system clock 5-66
- setting the time zone 5-66
- setting up a modem 4-19
- setting up a terminal 4-19
- setting up peripherals 4-12
- setting up printers 4-154
- setting up the line printer spooler 4-154
- shared code 2-65
- shared segment 2-64
- shell 3-25, 5-8
- shutdown
  - automatic unmounting of file systems 4-146
  - command syntax 3-28
  - `/etc/rcflag` 3-20
  - if system shutdown improperly A-1
  - improper 3-16
  - shutting down a cnode in an HP-UX cluster 3-31
  - shutting down an HP-UX cluster 3-31
  - shutting down the system 3-28, 5-1
  - system maintenance 3-28
- shutting down the recovery system 4-135
- slave pseudo device 4-36
- speed indicator 4-34
- spool 4-183
- SRM 4-70
- stack segment 2-59-60, 2-64
- stack size D-96
- standard error 2-9
- standard input 2-9
- standard output 2-9
- sticky bit 2-38, 2-66, 4-145
- summarizing and reporting accounting information 7-12, 7-51
- summary information area 2-16
- superblock A-4, 2-14, 2-16, 2-23, 2-26, 4-109, 4-140
- superblock, alternate A-8, 2-16
- superuser 2-10, 2-12
- suspending a user 5-20

- swap device 4-74, 4-77, 4-134, 6-6, 6-14
- swap space
  - applications 6-17
  - changing 4-129, 6-15, 6-26, 6-28
  - configuring 6-4, 6-13
  - description 2-58, 2-63, 6-13
  - determining requirements 6-16
  - enabling 4-76
  - enabling swap devices 6-7, 6-24
  - entry in configuration description file 6-22
  - HP-UX cluster 2-95, 6-56
  - multiple swap devices 6-13
  - operating system parameters D-96
  - requirements 4-109
- switch settings 4-13
- symbolic mode 2-34
- SYSBACKUP 2-15, 6-6
- SYSDEBUG 2-15, 3-14
- SYSHPUX 2-15, 3-14, 6-6
- system accounting
  - corrupted files 7-69
  - daily usage 7-5, 7-58
  - directory structure 7-15
  - errors 7-61
  - files 7-77
  - holidays 7-68
  - HP-UX cluster 2-76, 2-105, 7-4, 7-17, 7-19, 7-22
  - installation 7-3, 7-9
  - logging in 7-15
  - monthly usage 7-67
  - overview 7-1-2, 7-7
  - shell scripts 7-71
  - system data flow 7-13
- system administration mode 3-26, 3-28, 4-125, 4-127, 5-51
- system administrator duties 1-1, 1-6
- system backup on an HP-UX cluster 5-21, 5-26
- system calls
  - exec 2-6-7, 2-65, 2-67
  - fork 2-6
  - vfork 2-6
- system clock 5-22, 5-66

- system clock on an HP-UX cluster 5-66
- system console 3-9-10, 4-19
- system crash A-1, 5-21, 5-44
- system deadlock 2-40
- system maintenance 3-29
- system news 4-5
- system panic message 3-32
- system recovery 3-32
- system security 4-125, 4-132, 5-3, 5-14, 5-16, 5-50, 8-4
- system shutdown 3-16, 3-20, 3-28, 5-1
- system startup 3-1-2, 3-9, 3-12
- System V IPC 6-31

## T

### TABLES

- Access Permissions B-7
- Additional Settings for Handshake Type Bits 3-11
- Advantages/Disadvantages of Networked Workstations 2-70
- Advantages/Disadvantages of Timeshared Systems 2-69
- Application Execution Environment Functions C-4
- AXE Filesets and Their Dependencies C-9
- Choosing Port Usage 4-27
- Cluster Information Sheet 4-79
- D-1. Virtual Memory Parameter Table D-97
- Default Special Device Files 4-19
- Example Free Block Bitmap in an 8192/1024 File System 2-21
- Explanation of File Permission Bits of 6754 2-37
- Flexible Disk Label Colors C-3
- General Template for Ports 4-30
- General Template for Pseudo Devices 4-37
- HP-UX Cluster O.S. Parameters 6-53
- Kernel Driver & Major Numbers for Hard Disk Drives 4-45
- Kernel Drivers and Major Device Numbers 4-66
- Kernel Drivers and Major Device Numbers for Data Communications 4-70
- Kernel Drivers for Printers 4-166
- Line Control Switch Pack Settings 3-11
- Magnetic Tape Operation Bit Settings 4-64
- mkfs Defaults 4-111
- Optional Kernel Drivers 6-12
- PATH Variable B-12
- PE Filesets and Their Dependencies C-12

- Printer Sub-Addresses 4-164
- Programming Environment Functions C-7
- Remote Spooler Options 4-168
- Setting the Volume Number for Printers with CIPER Protocol 4-183
- Tape Density and Unit Number Bit Settings 4-64
- Volume Number with AMIGO Protocol or No Protocol 4-183
- tar** 5-72
- tcio** 5-71
- TERM environment variable 4-71
- TERM environment variable 3-25, 4-29
- terminal 4-30
  - adding 5-3
  - baud rate 4-34
  - non-HP 4-150, 4-189
  - problems B-9
  - removing 4-35, 5-20
  - setting up 4-19, 5-11
- terminfo database 4-71
- text size D-96
- time 5-66
- time zone on an HP-UX cluster 5-67
- time zone (TZ) 4-10-11, 5-66
- time-out option 4-33
- total accounting records 7-8, 7-21, 7-26, 7-48, 7-51
- transferring files 5-70
- trusted computing base 8-9
- trusted system
  - auditing 8-11
  - conversion 8-7
  - general 8-1, 8-37
  - networking 8-2
- tset** command 4-73
- turning off the computer 3-28
- TZ environment variable 3-25, 4-5, 4-10, 5-66

## **U**

- Unable to access ttyxx 4-35
- unattended mode 3-11-12, 3-14
- unmounting a file system 4-140, 4-144
- update
  - from cartridge tape 4-196

- from flexible disk 4-212
- information to check for 4-211, 4-226
- installing optional software 4-186
- non-HP terminals 4-189
- on an HP-UX cluster 4-188, 4-195
- overview 4-187, 4-189
- updating the system 4-3, 4-186, 4-191, 4-197, 5-27
- updating the HFS file system 2-23
- use count 2-65
- user ID 2-35
- user name 5-7
- user process size limits 6-31
- user survey form (see evaluating users' needs) 1-6
- users
  - adding 5-3, 5-12
  - adding to a group 5-13
  - communicating with 5-56
  - group 5-54
  - removing 5-3, 5-20
  - suspending 5-20
- using the recovery system manually 4-135
- /usr/lib/terminfo 4-8
- /usr/lib/terminfo database 4-71
- UUCP 5-76
- UUCP in an HP-UX cluster 2-105

## **V**

- virtual memory parameters D-97
- virtual terminal (VT) 5-77
- volume directory information 2-15
- volume header 2-14
- volume number
  - magnetic tape drive 4-64
  - terminals and modems 4-31

## **W**

- wall 5-56
- write command 5-56
- write permission 2-32, 2-35
- write protection 2-46
- write-protect mechanism

cartridge tape 4-196  
flexible disk 4-212







**HP Part Number**  
**98594-90060**

Microfiche No. 98594-99060  
Printed in U.S.A. E1288



**98594-90660**

For Internal Use Only