# CONCENTRIX-2800

# System Administration

January 1992

This manual provides guidelines and alphabetic references to commands and procedures used in the administration of the Concentrix-2800 operating system.


Part number: 301-00500
Revision D

Model Number: D002.18

Use the model number for ordering purposes. The model number includes the text of the document and any ancillary parts such as binder, front and rear covers, and label. The part number refers to just this part - the text of the document.

Concentrix-2800 Version: 2.2.0



ALLIANT COMPUTER SYSTEMS CORPORATION
One Monarch Drive ■ Littleton ■ Massachusetts 01460
508-486-4950

Alliant Computer Systems Corporation reserves the right to make changes in specifications and other information contained in this document without prior notice.

Although due care has been taken to present accurate information, ALLIANT DISCLAIMS ALL WARRANTIES WITH RESPECT TO THE CONTENTS OF THIS DOCUMENT (INCLUDING WITHOUT LIMITATION WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE), EITHER EXPRESSED OR IMPLIED. ALLIANT SHALL NOT BE LIABLE FOR DAMAGES RESULTING FROM ANY ERROR CONTAINED HEREIN, INCLUDING, BUT NOT LIMITED TO, FOR ANY SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF, OR IN CONNECTION WITH, THE USE OF THIS DOCUMENT.

**WARNINGS:** The software described in this document runs on the Alliant FX/2800 and FX/800 Series. This equipment generates, uses, and can radiate radio frequency energy, and, if not installed and used in accordance with instructions in the manual, may cause interference with radio communications.

It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, and with the Radio Interference Regulations, CRC c. 1374. These rules are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference, in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

*Le present appareil numerique n'emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de classes A prescrites dans le reglement sur le brouillage radioelectrique edicte par le ministere des communications du Canada.*

The information contained in this document is summary in nature. More detailed information is available from Alliant.

This software product and its documentation set are copyrighted (unpublished) and all rights are reserved by Alliant Computer Systems Corporation. Usage of this product is only allowed under the terms set forth in the Alliant Purchase and License Agreement. Any reproduction or distribution of this document, in whole or in part, without the prior written consent of Alliant Computer Systems Corporation is prohibited.

Concentrix and Concentrix-2800 software and documentation is based in part on the Fourth Berkeley Software Distsribution under license from the The Regents of the University of California.

GNU Emacs software and documentation is based on the *GNU Emacs Manual* by Richard Stallman, 6th edition, version 18, March, 1987, distributed by the Free Software Foundation.

Copyright © 1992
ALLIANT COMPUTER SYSTEMS CORPORATION
All Rights Reserved

## Revision History

| | Date | Concentrix-2800 Release | Comments |
|---|---|---|---|
| A | July 1990 | V1.0 | initial version |
| B | Dec 1990 | V2.0 | update |
| C | Mar 1991 | V2.1 | update |
| D | Jan 1992 | V2.2 | update |

Trademarks of Alliant Computer Systems Corporation:

| | | | |
|---|---|---|---|
| Alliant | FX/1 | FX/Series | FX/FORTRAN |
| Concentrix | FX/4 | FX/40 | FX/C |
| Diagnostix | FX/8 | FX/80 | FX/Ada |
| DNX | VIP | FX/80T | Visualization Series |
| ANSR | ANSR/LCX | FX/82 | FX/RT |
| VFX/4 | VFX/40 | VFX/80 | VFX/82 |
| Visual Supercomputer | DNX-2800 | SRM/1 | Visual Supercomputing |
| FX/2800 Series | FX/2800 | Concentrix-2800 | Scalar/C |
| FX/C-2800 | FX/FORTRAN-2800 | FX/RT-2800 | Diagnostix-2800 |
| FX/Ada-2800 | FASTf | FX/800 | CAMPUS |
| | CAMPUS/800 | Clusternode | |
| Raster Technologies | GX4000 | RNeWS | |
| Model One | Model One/10 | Model One/25 | Model One/40 |
| Model One/60 | Model One/75 | Model One/85 | Model One/380 |
| ONELIB | ONELIBC | ONELIBF77 | Model One/385 |

Other trademarks that may appear in this document include:

| | |
|---|---|
| CCA EMACS | Trademark of Uniworks, Inc. |
| NeWS, NFS,NIS<br>Sun-3, Sun-4, Sun-3/260<br>Sun-4/260 | Trademarks of Sun Microsystems, Inc. |
| UNIX | Registered trademark of AT&T in the United States and Other Countries |
| Multibus, i860 | Trademarks of Intel Corporation |
| DEC,DECnet, VAX,<br>Digital Command Language (DCL),<br>VAX/VMS, DNA,<br>VAXcluster | Trademarks of Digital Equipment Corporation |
| Cray-1, Cray X-MP,<br>Cray-2 | Trademark of Cray Research, Inc. |
| Ethernet | Trademark of Xerox Corporation |
| Ada | Registered trademark of the U.S. government, Ada Joint Program Office (AJPO) |
| Network Computing System,<br>NCS | Registered trademarks of Apollo Computer, Inc. |
| Hyperchannel | Trademark of Network Systems Corp. |
| X11,<br>X-Window System | Trademark of Massachusetts Institute of Technology |
| Postscript | Registered trademark of Adobe Systems |
| Para-Linpack/FX,<br>Para-Eispack/FX | Trademarks of Kuck and Associates, Inc. |
| CommUnity | Registered trademark of Technology Concepts, Inc. |
| UltraNet | Trademark of Ultra Network Technologies Corp. |
| VCL, EDT+ | Boston Business Computing, Ltd. |
| KiNET, DEKnet | Trademarks of Ki Research, Inc. |
| EXPRESS | Trademark of ParaSoft Corporation |

Printed in the United States of America

# TABLE OF CONTENTS


PREFACE

CHAPTER 1 BASIC OPERATIONS

1.1 FRONT PANEL AND CONSOLE TERMINAL ..... 1-3
1.2 HARDWARE CONFIGURATION ..... 1-4
1.2.1 Processor Module ..... 1-4
1.2.2 Input/Output Module ..... 1-5
1.2.3 Computational Configurations ..... 1-6
1.2.4 Memory Subsystem ..... 1-7
1.3 SOFTWARE CONFIGURATION ..... 1-8
1.3.1 Users and Superusers ..... 1-8
1.3.2 File System ..... 1-8
1.3.2.1 Directories Provided with Concentrix-2800 ..... 1-9
1.3.2.2 Files that Grow ..... 1-10
1.3.2.3 Files that Need Maintenance ..... 1-11
1.3.2.4 Listing Files and Directories ..... 1-12
1.3.3 Modes of Operation ..... 1-13
1.3.3.1 Console Mode ..... 1-14
1.3.3.2 VMEbus Booter Mode ..... 1-16
1.3.3.3 Alliant ddt Mode ..... 1-17
1.3.3.4 Diagnostix-2800 Mode ..... 1-17
1.3.3.5 FX/2800 Startup Mode ..... 1-18
1.3.3.6 FX/2800 Booter Mode ..... 1-20
1.3.3.7 Alliant ddt 2800 Mode ..... 1-21
1.3.3.8 Single-User Mode ..... 1-21
1.3.3.9 Concentrix-2800 Mode ..... 1-22
1.4 START-UP AND SHUTDOWN ..... 1-22
1.4.1 Shutdown Versus Reboot ..... 1-23
1.4.2 Fastboot ..... 1-23
1.4.3 Boot Device ..... 1-23
1.4.3.1 Booter Program ..... 1-23
1.4.3.2 Diagnostix-2800 Root Partition ..... 1-24
1.4.3.3 Concentrix-2800 Root Partition ..... 1-25
1.4.4 Boot Script ..... 1-25
1.4.5 System Initialization ..... 1-31
1.4.5.1 Script rc.diag ..... 1-31
1.4.5.2 Script rc ..... 1-32
1.4.5.3 Script rc.local ..... 1-34
1.4.6 Boot from RAM Disk ..... 1-37
1.5 DATE AND TIME ..... 1-38
1.6 SYSTEM FAILURES ..... 1-39
1.6.1 Terminal Hang ..... 1-39
1.6.2 System Hang ..... 1-40
1.6.3 System Panic ..... 1-40
1.6.4 Peripheral Errors ..... 1-41

## CHAPTER 2 USER ACCOUNTS


2.1 LOGIN PROCEDURE ..... 2-3
2.2 PASSWD FILE ..... 2-4
2.3 GROUP FILE ..... 2-5
2.4 SKELETON FILES ..... 2-6
2.5 USER COMMUNICATIONS ..... 2-7
2.6 MAIL ..... 2-7
2.6.1 Mail Command ..... 2-7
2.6.2 Mail Queue (User Mailboxes) ..... 2-7
2.6.3 Start-up Files ..... 2-8
2.6.4 Alias Data Base ..... 2-8
2.6.5 Encrypted Mail ..... 2-8
2.6.6 Sendmail Configuration ..... 2-9
2.7 ACCOUNTING ..... 2-10
2.7.1 Standard Resource Accounting ..... 2-10
2.7.2 Connect-time Accounting ..... 2-11
2.7.3 Custom Resource Accounting ..... 2-11
2.8 FILE SYSTEM DISK QUOTAS ..... 2-13


## CHAPTER 3 DEVICES


3.1 SPECIAL FILES ..... 3-2
3.1.1 Making Special Files ..... 3-4
3.1.2 Listing Special Files ..... 3-5
3.1.3 Concentrix-2800 Supported Devices ..... 3-5
3.1.4 Input/output Address Ranges ..... 3-7
3.1.5 Diagnostix-2800 Supported Devices ..... 3-8
3.1.6 Default Special Files ..... 3-10
3.2 DISK DEVICES ..... 3-11
3.2.1 Disk Names ..... 3-11
3.2.2 Supported Disk Drives ..... 3-11
3.2.3 Disk Geometry ..... 3-12
3.2.3.1 Disk Storage Layout ..... 3-12
3.2.3.2 Disk Labels ..... 3-13
3.2.3.3 Partitions ..... 3-13
3.2.3.4 Flaws ..... 3-16
3.2.4 Disk Layout Considerations ..... 3-17
3.2.4.1 Protecting Disks ..... 3-18
3.2.4.2 Unmarking Partitions and Removing Unused Partitions ..... 3-18
3.2.4.3 Swapping and Paging Considerations ..... 3-19
3.2.4.4 Non-Standard User Partitions ..... 3-20
3.3 SMALL COMPUTER SYSTEM INTERFACE (SCSI) DEVICES 3-21
3.3.1 SCSI Disk Devices ..... 3-21
3.3.1.1 Supported SCSI Disk Drives ..... 3-22
3.3.1.2 SCSI Disk Labels ..... 3-22
3.3.1.3 SCSI Flaws ..... 3-23
3.3.2 SCSI Tape Devices ..... 3-24
3.3.3 SCSI Switch Configuration ..... 3-24
3.3.4 Loading an 8mm Cartridge ..... 3-25
3.3.5 Unloading an 8mm Cartridge ..... 3-26
3.3.6 8mm Cartridge Power-up Initialization Indicators ..... 3-26
3.4 TAPE DEVICES (9-TRACK) ..... 3-27
3.4.1 Tape Names ..... 3-27

| | | |
|---|---|---|
| 3.4.2 | Supported 9-Track Tape Drives | 3-27 |
| 3.4.3 | Configuring Tape Drives | 3-28 |
| 3.4.4 | Tape Operations | 3-29 |
| 3.4.4.1 | Loading a Telex Tape | 3-29 |
| 3.4.4.2 | Unloading a Telex Tape | 3-29 |
| 3.4.4.3 | Loading a Cipher Tape | 3-30 |
| 3.4.4.4 | Unloading a Cipher Tape | 3-30 |
| 3.4.4.5 | Examining a Tape | 3-31 |
| 3.4.5 | Special Tape Drives | 3-32 |
| 3.4.5.1 | Modifying a Structure for a Special Tape Drive | 3-32 |
| 3.4.5.2 | Configuring Special Tape Drives | 3-34 |
| 3.5 | TERMINAL DEVICES | 3-35 |
| 3.5.1 | Terminal Names | 3-35 |
| 3.5.1.1 | Script mknewdevs | 3-36 |
| 3.5.1.2 | Command mkttynod | 3-36 |
| 3.5.1.3 | Command mknod | 3-36 |
| 3.5.1.4 | Pseudo Terminals | 3-37 |
| 3.5.2 | Hardware Configuration | 3-37 |
| 3.5.2.1 | Cabling | 3-38 |
| 3.5.2.2 | Switches | 3-40 |
| 3.5.3 | Software Configuration | 3-42 |
| 3.5.3.1 | Logging In from An Apollo Window | 3-43 |
| 3.6 | SYSTEM MEMORY MAPS | 3-44 |

**CHAPTER 4 FILE SYSTEMS**

| | | |
|---|---|---|
| 4.1 | TYPES OF FILE SYSTEMS | 4-4 |
| 4.2 | MAKING NEW FILE SYSTEMS | 4-4 |
| 4.3 | DUMPING, RESTORING, AND MOVING FILE SYSTEMS | 4-7 |
| 4.3.1 | Dump Levels | 4-7 |
| 4.3.2 | Dump Procedure | 4-8 |
| 4.3.3 | Dump Media | 4-10 |
| 4.3.4 | Restoring a Full File System | 4-11 |
| 4.3.5 | Restoring Files Selectively to a File System | 4-12 |
| 4.3.6 | Moving a File System | 4-13 |
| 4.3.7 | Saving and Moving Directories with bfio | 4-14 |
| 4.3.8 | Saving and Moving Directories with tar | 4-15 |
| 4.3.9 | Restoring The Root Partition | 4-16 |
| 4.3.9.1 | Restoring the Diagnostix-2800 Root Partition | 4-16 |
| 4.3.9.2 | Restoring the Concentrix-2800 Root Partition | 4-18 |
| 4.3.10 | Cleaning Up /tmp | 4-19 |
| 4.4 | TUNING AND FIXING | 4-21 |
| 4.4.1 | Examining a File System | 4-21 |
| 4.4.2 | File System Checks | 4-22 |
| 4.4.3 | File System Layout | 4-23 |
| 4.4.3.1 | Cylinder Groups | 4-23 |
| 4.4.3.2 | Inodes, Blocks, and Fragments | 4-24 |
| 4.4.3.3 | Free Space | 4-25 |
| 4.4.3.4 | Rotational Delay | 4-26 |
| 4.4.3.5 | Superblock | 4-26 |

## CHAPTER 5 MONITORING AND TUNING


| | | |
|---|---|---|
| 5.1 | ACTIVE SYSTEM | 5-3 |
| 5.1.1 | System Configuration and Current Users | 5-3 |
| 5.1.2 | Monitor Display | 5-4 |
| 5.1.3 | Process Status | 5-6 |
| 5.1.4 | Process Killing (Hung Process) | 5-8 |
| 5.2 | FILE SYSTEM | 5-9 |
| 5.3 | SCHEDULING COMPUTATIONAL RESOURCES | 5-11 |
| 5.3.1 | Computational Resource Types | 5-11 |
| 5.3.2 | Execution Modes | 5-12 |
| 5.3.2 | Scheduling Classes | 5-13 |
| 5.3.3 | Scheduling Vectors | 5-14 |
| 5.3.4 | Computational Clusters | 5-16 |
| 5.3.4.1 | Cluster Dynamics | 5-17 |
| 5.3.4.2 | Startup Cluster Configuration (console.config) | 5-18 |
| 5.3.4.3 | Dynamic Cluster Reconfiguration | 5-20 |
| 5.3.4.4 | Cluster Configuration Startup File (/etc/rc.cluster) | 5-21 |
| 5.3.5 | Modifying Scheduling Vectors | 5-22 |
| 5.3.5.1 | High Concurrent Use | 5-23 |
| 5.3.5.2 | High Detached Use | 5-25 |
| 5.3.5.3 | New Clusters | 5-25 |
| 5.3.5.4 | Clusters of Different Sizes | 5-26 |
| 5.3.5.5 | High (Low) Priority Work | 5-27 |
| 5.3.5.6 | Dedicated Resources | 5-28 |
| 5.3.5.7 | Implicit Migration to IP | 5-29 |
| 5.3.6 | Real-time Scheduler | 5-29 |
| 5.4 | CRONTAB | 5-30 |
| 5.5 | KERNEL VARIABLES | 5-31 |


## CHAPTER 6 PRINT AND BATCH QUEUES


| | | |
|---|---|---|
| 6.1 | PRINTERS AND PRINT QUEUES | 6-3 |
| 6.1.1 | User Capabilities | 6-4 |
| 6.1.2 | General Administration | 6-5 |
| 6.1.2.1 | Line Printer Control (lpc) Program | 6-5 |
| 6.1.2.2 | Trouble-shooting and Error Messages | 6-7 |
| 6.1.2.3 | Error Messsages (lpc) | 6-8 |
| 6.1.2.4 | Error Messsages (lpq) | 6-8 |
| 6.1.2.5 | Error Messsages (lpr) | 6-9 |
| 6.1.2.6 | Error Messsages (lprm) | 6-9 |
| 6.1.3 | Default Printer | 6-10 |
| 6.1.4 | Additional Parallel Printers | 6-11 |
| 6.1.5 | Serial Printers | 6-12 |
| 6.1.6 | Remote Printers | 6-13 |
| 6.1.7 | Output Filters | 6-14 |
| 6.1.8 | Printer Accounting | 6-15 |
| 6.2 | BATCH QUEUES | 6-16 |
| 6.2.1 | Batch Setup and Queue Configuration | 6-17 |
| 6.2.2 | User Capabilities and General Administration | 6-17 |

## CHAPTER 7 NETWORKS


| | | |
|---|---|---|
| 7.1 | LOCAL AREA NETWORKS | 7-5 |
| 7.1.1 | Internet Services | 7-6 |
| 7.1.2 | Device Names and Addresses | 7-8 |
| 7.1.2.1 | Ethernet Controller Names | 7-8 |
| 7.1.2.2 | Ethernet Addresses | 7-9 |
| 7.1.2.3 | Pseudo Terminals | 7-9 |
| 7.1.3 | Network Names and Addresses | 7-10 |
| 7.1.3.1 | Setting localhost for non-networked Machine | 7-12 |
| 7.1.4 | Equivalent (Trusted) Systems And Accounts | 7-13 |
| 7.1.5 | Examining and Testing the Network | 7-14 |
| 7.1.6 | Ethernet Performance | 7-15 |
| 7.1.6.1 | Placement of Softnet Portion of Ethernet Driver | 7-15 |
| 7.1.6.2 | Buffers | 7-15 |
| 7.2 | UUCP NETWORK | 7-17 |
| 7.3 | NETWORK INFORMATION SERVICE | 7-19 |


## APPENDIX A ADMINISTRATION COMMANDS


| | | |
|---|---|---|
| ac | login accounting | A-2 |
| accton | enable/disable system accounting | A-3 |
| arcv | convert archives to new format (old) | A-4 |
| arp | address resolution display and control | A-5 |
| automount | automatically mount NFS file systems | A-7 |
| batchd | batch daemon | A-12 |
| biod | NFS daemon | A-14 |
| bootparamd | boot parameter server | A-15 |
| catman | create on-line documentation output directory | A-16 |
| chown | change owner | A-18 |
| clcr | create cluster | A-19 |
| clinfo | print cluster information | A-21 |
| clri | clear i-node | A-22 |
| clrm | remove cluster | A-23 |
| compile | compile the terminfo database | A-24 |
| comsat | biff server | A-25 |
| create_root | generate root partition for diskless client | A-26 |
| cron | clock daemon | A-27 |
| dfmon | dynamic disk summary and usage | A-29 |
| diskf | disk utility program | A-30 |
| dlabel | generate disk labels from prototype file | A-38 |
| dmesg | collect system diagnostic messages in error log (old) | A-40 |
| dump | incremental file system backup | A-41 |
| dumpfs | dump file system information | A-45 |
| dumpterm | print contents of compiled terminfo file in human-readable form | A-49 |
| edquota | edit user quotas | A-50 |
| exportfs | export and unexport directories to NFS clients | A-52 |
| extracting | extract from a tar formatted software distribution tape | A-54 |
| extendswap | enable/disable multiple-disk swapping/paging area | A-55 |
| fastboot | reboot system without checking disks | A-56 |
| fasthalt | halt system without checking disks | A-58 |
| fingerd | remote user information server | A-59 |

| | | |
|---|---|---|
| fix_bootparams | add or delete entries from /etc/bootparams | A-60 |
| fsck | file system consistency check / interactive repair | A-61 |
| fsirand | install random i-node generation numbers | A-81 |
| ftpd | DARPA Internet file transfer protocol server | A-82 |
| gettable | get NIC format host tables from a host | A-85 |
| getty | set terminal mode | A-86 |
| halt | halt system | A-88 |
| htable | convert NIC standard format host tables | A-89 |
| ifconfig | configure network interface parameters | A-90 |
| ifstats | print out ethernet board statistics | A-93 |
| implog | IMP log interpreter | A-95 |
| implogd | IMP logger process | A-96 |
| inetd | internet services daemon | A-97 |
| init | system initialization and process control | A-100 |
| INSTALL | install diskless/NFS architecture and clients | A-102 |
| keyenvoy | talk to keyserver | A-106 |
| keyserv | server for storing public and private keys | A-107 |
| kgmon | control kernel profiling | A-108 |
| kgprof | display kernel execution profile | A-110 |
| lockd | network lock daemon | A-113 |
| lpc | line printer control program | A-114 |
| lpd | line printer daemon | A-116 |
| mailq | print contents of queue | A-119 |
| mailstats | print sendmail statistics | A-120 |
| makedbm | make a network information service dbm file | A-121 |
| make_dirs | make a directory and its parents | A-123 |
| makefs | construct a file system | A-124 |
| MAKEHOSTS | set up /usr/hosts directory | A-127 |
| makekey | generate encryption key | A-128 |
| makeramdisk | create RAM disk image | A-129 |
| makewhatis | construct whatis data base | A-132 |
| mconnect | connect to SMTP mail server socket | A-133 |
| mkdisknod | make special files for standard and scsi disk devices | A-134 |
| mkfile | create NFS swap file | A-135 |
| mkhosts | generate hashed host table | A-136 |
| mklost+found | make lost+found directory for fsck | A-137 |
| mknewdevs | create default devices for initial installation | A-138 |
| mknod | build special file or make a named pipe | A-142 |
| mkpasswd | generate hashed password table | A-143 |
| mkproto | construct a prototype file system | A-144 |
| mkptynod | make special files for pseudo terminals | A-146 |
| mkstripe | make striped partition | A-147 |
| mkswap | make node for swapping partition | A-148 |
| mktape | make special files for tapes | A-149 |
| mkttynod | make special files for terminals | A-150 |
| mount | mount file systems | A-151 |
| mountd | NFS mount request server | A-155 |
| named | Internet domain name server | A-156 |
| ndbootd | ND boot block server | A-161 |
| netdae | batch network server daemon | A-162 |
| netrecv | batch network receive daemon | A-163 |

| | | |
|---|---|---|
| netsend | batch network send daemon | A-164 |
| newaliases | rebuild data base for mail aliases file | A-165 |
| newkey | create a new key in the publickey database | A-166 |
| newfs | construct a new file system (unsupported) | A-167 |
| nfsd | NFS daemon | A-168 |
| nfsstat | NFS statistics | A-169 |
| nslookup | query Internet name servers interactively | A-171 |
| ntalkd | server for talk program | A-176 |
| opt_software | select software from a distribution tape | A-177 |
| pac | printer/plotter accounting information | A-178 |
| patchunix | patch Concentrix-2800 image | A-179 |
| ping | send icmp echo packet (network debugging) | A-182 |
| portmap | DARPA port to RPC program number mapper | A-183 |
| ptynum | change nptys flag in kernel image | A-184 |
| quot | summarize file system ownership | A-185 |
| quotacheck | file system quota consistency checker | A-186 |
| quotaon quotaoff | turn file system quotas on and off | A-187 |
| rc | command scripts for starting system | A-189 |
| rdate | set system data from a remote host | A-191 |
| rdump | file system dump across the network | A-192 |
| reboot | reboot system | A-193 |
| renice | alter priority of running processes | A-195 |
| reportcfg | report software/hardware configuration | A-196 |
| repquota | summarize quotas for a file system | A-197 |
| restore | incremental file system restore | A-198 |
| rexd | RPC-based remote execution server | A-203 |
| rexecd | remote execution server | A-204 |
| rlogind | remote login server | A-206 |
| rmail | handle remote mail received via uucp | A-209 |
| rmt | remote magtape protocol module | A-210 |
| route | manually manipulate the routing tables | A-212 |
| routed | network routing daemon | A-214 |
| rpcinfo | report RPC information | A-217 |
| rquotad | remote quota server | A-219 |
| rrestore | restore a file system dump across the network | A-220 |
| rshd | remote shell server | A-221 |
| rstatd | kernel statistics server | A-223 |
| rtnamed | logical name daemon | A-224 |
| rtncrtbl | create logical name table | A-226 |
| rtndeltbl | delete logical name table | A-228 |
| rtndumptbl | dump logical name table | A-229 |
| rtnenter | enter logical name into table | A-230 |
| rtnlookup | look up logical name in table | A-232 |
| rtnpurge | purge object from logical name tables | A-233 |
| rtnremove | remove logical name from table | A-234 |
| rusersd | network username server | A-235 |
| rwalld | network rwall server | A-236 |
| rwhod | system status server | A-237 |
| sa | maintain system accounting files | A-239 |
| sendmail | send mail over the internet | A-241 |
| setclass | set process class | A-246 |
| setclassdef | change setclassdef flag in kernel image | A-248 |

| | | |
|---|---|---|
| setcomplex | set or display thresholds for computational complex | A-250 |
| setcres | set computing resource | A-252 |
| setdefclass | set default scheduling class | A-254 |
| setmaxshare | change maxsharesize flag in kernel image | A-256 |
| setpn | set the maximum number of processors for a concurrent job | A-258 |
| setsched | set scheduling vector for computing resource | A-259 |
| settype | set computing resource type | A-263 |
| setup_client | add or remove diskless client to/from server | A-264 |
| setup_exec | install diskless client software on a server | A-266 |
| showdefclass | display default scheduling classes | A-267 |
| showmount | show all remote mounts | A-269 |
| showsched | display scheduling vectors | A-270 |
| shutdown | shut down system at a given time | A-271 |
| software | install layered product | A-274 |
| spray | spray packets | A-277 |
| sprayd | spray server | A-278 |
| srboot | Alliant FX/2800 booter | A-279 |
| srcreate | create shared memory region | A-281 |
| srdelete | delete shared memory region | A-282 |
| srdisplay | display shared memory regions | A-283 |
| statd | network status monitor | A-284 |
| sticky | persistent text and append-only directories | A-285 |
| sunboot | start Sun system | A-286 |
| swapon | specify additional device for paging and swapping (old) | A-289 |
| sync | update the super block | A-290 |
| syslog | log system messages (old) | A-291 |
| syslogd | log system messages | A-292 |
| telnetd | DARPA TELNET protocol server | A-294 |
| tftpd | DARPA trivial file transfer protocol server | A-295 |
| tnamed | DARPA trivial name server | A-296 |
| timed | time server daemon | A-297 |
| timedc | timed control program | A-299 |
| timezone | set time zone | A-301 |
| trpt | transliterate protocol trace | A-302 |
| trsp | transliterate sequenced packet protocol trace | A-304 |
| tunefs | tune up an existing file system | A-306 |
| umount | unmount file systems | A-308 |
| unmkstripe | disassemble disk stripe | A-310 |
| update | periodically update the superblock | A-311 |
| uucico | transfer files queued by uucp or uux | A-312 |
| uuclean | UUCP spool directory clean-up | A-315 |
| uucpd | UUCP daemon | A-316 |
| uupoll | poll a remote UUCP site | A-317 |
| uusnap | show snapshot of the UUCP system | A-318 |
| uuxqt | UUCP execution file interpreter | A-319 |
| uv | print kernel version information | A-320 |
| verify_tapevol_arch | verify mounted distribution tape is correct | A-321 |
| vipw | edit the password file | A-322 |
| vmlock | enable/disable process swapping and outpaging | A-323 |
| xdrtoc | convert distribution table of contents | A-324 |
| xdvstat | xdv disk statistics dynamic display | A-325 |
| ypbind | network information service binder process | A-328 |

ypinit build and install network information service data base ..... A-330
ypmake rebuild network information sevice data base ..... A-331
yppasswdd server for modifying network information service password file A-333
yppoll network information service get map version ..... A-334
yppull network information service get map from master ..... A-335
yppush force propagation of a changed NIS map ..... A-336
ypserv network information service server process ..... A-338
ypset point ypbind at a particular server ..... A-340
ypupdated server for changing NIS information ..... A-341
ypwhich determine network information service master ..... A-342
ypxfr transfer an NIS map ..... A-343
zic timezone compiler ..... A-345

APPENDIX B ADMINISTRATION PROCEDURES

Add_ethernet add system to Ethernet ..... B-2
Add_printer add printer to system ..... B-6
Add_quotas add file system quotas ..... B-8
Add_terminal add terminal to system ..... B-10
Add_user add a new user to the system ..... B-13
Add_uucp add system to UUCP network ..... B-17
Add_yp add network information service (NIS) to network ..... B-20
Boot_system start system and bring to multi-user mode ..... B-23
Broadcast broadcast messages to users ..... B-34
Build_kernel build custom kernel ..... B-36
Call_cs report problem to Alliant customer service ..... B-37
Connect_diag connect to Diagnostix-2800 ..... B-38
Disable_terminal
disable (enable) interactive terminal ..... B-39
Extend_swap extend area for swapping and paging ..... B-40
Format_disk format and label disk ..... B-42
Halt_system halt system in console mode ..... B-45
Install_batch install batch system ..... B-46
Install_mail install mail on a new system ..... B-49
Label_disk label (relabel) disk ..... B-51
Make_file_system
make and mount a new file system ..... B-58
Patch_kernel modify variables in the system kernel ..... B-60
Recover_system
recover from hang or panic ..... B-66
Remove_user remove a user from the system ..... B-70
Shutdown_system
shut down system to single-user mode ..... B-73

APPENDIX C SENDMAIL

C.1 DESIGN GOALS ..... C-2
C.2 OVERVIEW ..... C-3
C.2.1 System Organization ..... C-3
C.2.2 Interfaces to the Outside World ..... C-4
C.2.2.1 Argument Vector/exit Status ..... C-4
C.2.2.2 SMTP Over Pipes ..... C-4
C.2.2.3 SMTP Over an IPC Connection ..... C-4
C.2.3 Operational Description ..... C-4

| | | |
|---|---|---|
| C.2.3.1 | Argument Processing and Address Parsing | C-4 |
| C.2.3.2 | Message Collection | C-5 |
| C.2.3.3 | Message Delivery | C-5 |
| C.2.3.4 | Queueing for Retransmission | C-5 |
| C.2.3.5 | Return to Sender | C-5 |
| C.2.4 | Message Header Editing | C-6 |
| C.2.5 | Configuration File | C-6 |
| C.3 | USAGE AND IMPLEMENTATION | C-6 |
| C.3.1 | Arguments | C-6 |
| C.3.2 | Mail to Files and Programs | C-7 |
| C.3.3 | Aliasing, Forwarding, Inclusion | C-7 |
| C.3.3.1 | Aliasing | C-7 |
| C.3.3.2 | Forwarding | C-7 |
| C.3.3.3 | Inclusion | C-8 |
| C.3.4 | Message Collection | C-8 |
| C.3.5 | Message Delivery | C-8 |
| C.3.6 | Queued Messages | C-9 |
| C.3.7 | Configuration | C-9 |
| C.3.7.1 | Macros | C-9 |
| C.3.7.2 | Header Declarations | C-9 |
| C.3.7.3 | Mailer Declarations | C-10 |
| C.3.7.4 | Address Rewriting Rules | C-10 |
| C.3.7.5 | Option Setting | C-10 |
| C.4 | COMPARISON WITH OTHER MAILERS | C-10 |
| C.4.1 | Delivermail | C-10 |
| C.4.2 | MMDF | C-11 |
| C.4.3 | Message Processing Module | C-11 |
| C.5 | EVALUATIONS AND FUTURE PLANS | C-12 |
| C.6 | REFERENCES | C-13 |
| C.7 | SENDMAIL INSTALLATION AND OPERATION | C-15 |
| C.7.1 | Basic Installation | C-15 |
| C.7.1.1 | Off-the-shelf Configurations | C-16 |
| C.7.1.2 | Installation Using the Makefile | C-16 |
| C.7.1.3 | Installation by Hand | C-17 |
| C.7.2 | Normal Operations | C-20 |
| C.7.2.1 | The System Log | C-20 |
| C.7.2.2 | The Mail Queue | C-20 |
| C.7.2.3 | The Alias Database | C-23 |
| C.7.2.4 | Per-User Forwarding (.forward Files) | C-24 |
| C.7.2.5 | Special Header Lines | C-25 |
| C.7.3 | Arguments | C-25 |
| C.7.3.1 | Queue Interval | C-25 |
| C.7.3.2 | Daemon Mode | C-25 |
| C.7.3.3 | Forcing the Queue | C-25 |
| C.7.3.4 | Debugging | C-26 |
| C.7.3.5 | Trying a Different Configuration File | C-26 |
| C.7.3.6 | Changing the Values of Options | C-26 |
| C.7.4 | Tuning | C-27 |
| C.7.4.1 | Timeouts | C-27 |
| C.7.4.2 | Delivery Mode | C-28 |
| C.7.4.3 | Log Level | C-28 |
| C.7.4.4 | File Modes | C-29 |
| C.7.5 | Configuration File Details | C-30 |

C.7.5.1 The Syntax ..... C-30
C.7.5.2 The Semantics ..... C-33
C.7.5.3 Building a Configuration File From Scratch ..... C-38
C.7.6 Command Line Arguments ..... C-44
C.7.7 Configuration Options ..... C-45
C.7.8 Mailer Flags ..... C-46
C.7.9 Other Configuration Changes ..... C-48
C.7.9.1 Parameters in md/config.m4 ..... C-48
C.7.9.2 Parameters in src/conf.h ..... C-48
C.7.9.3 Configuration in src/conf.c ..... C-50
C.7.10 Summary of Support Files ..... C-52

APPENDIX D DISK MAINTENANCE AND INSTALL UTILITIES

D.1 CONTENTS OF RAM DISKS ..... D-1
D.1.1 Concentrix-2800 RAM Disk ..... D-1
D.1.2 Diagnostix-2800 Maintenance Utilities RAM Disk ..... D-2
D.1.3 Diagnostix-2800 Install Utilities RAM Disk ..... D-3
D.2 PROCEDURES FOR USING RAM DISK UTILITIES ..... D-4
D.2.1 Concentrix-2800 RAM Disk Utilities ..... D-4
D.2.2 Diagnostix-2800 RAM Disk Maintenance Utilities ..... D-6
D.2.3 Diagnostix-2800 RAM Disk Install Utilities ..... D-8
D.3 DEVICES ..... D-9
D.4 MOUNTING THE CONCENTRIX-2800 ROOT DISK ..... D-10
D.5 MOUNTING THE DIAGNOSTIX-2800 ROOT DISK ..... D-10

# PREFACE

This manual provides detailed guidelines and reference information on topics related to administration of the Concentrix-2800 operating system.

## AUDIENCE

This document is intended for people who administer a Concentrix-2800 operating system.

## STRUCTURE OF THE MANUAL

This manual consists of the following chapters and appendices. (Appendices A and B are also in Section 8 of the man on-line documentation database.)

| | |
|---|---|
| Chapter 1 | Introduces the system; explains how to start and stop the system. |
| Chapter 2 | Explains how to create and administer user accounts. |
| Chapter 3 | Explains how to create and administer devices. |
| Chapter 4 | Explains how to create and administer file systems. |
| Chapter 5 | Provides guidelines for monitoring and tuning the system. |
| Chapter 6 | Describes how to use printers, print queues, and batch queues. |
| Chapter 7 | Describes networks and UUCP administration. |
| Appendix A | Describes the administration commands in alphabetical order. |
| Appendix B | Describes various procedures for administering the system. These are procedures that you must step through, not automated scripts. |
| Appendix C | Provides information on the design of sendmail and the procedures for configuring the utility. This information is not needed unless you are modifying the configuration supplied by Alliant. |
| Appendix D | Describes the RAM disk maintenance and install utilities. |

## DIFFERENCES FROM THE PREVIOUS VERSION

Version D of *Concentrix-2800 System Administration* describes Version 2.2 of the Concentrix-2800 operating system. The information that has been added, changed, or deleted since the release of Version 2.1 of Concentrix-2800 is listed below.

### Sections Added/Changed

| | |
|---|---|
| Chapter 1 | Added information in section 1.2.4, "Memory Subsystem." |
| Chapter 2 | Added section 2.8, "Filesystem Disk Quotas." |
| Chapter 3 | Changed table on disk partition sizes, added FDDI and Berkeley Packet Filter devices, adjusted list of default SCSI devices. |
| Chapter 4 | Changes and additions to section 4.4.3 "File System Layout." |
| Chapter 5 | corrected information in section 5.3.2, "Execution Modes." |
| Chapter 7 | Section on setup for non-networked system added. |

### Commands Added in Appendix A

| | |
|---|---|
| automount | automatically mount NFS file system |
| dfmon | dynamic disk summary and usage |
| edquota | edit user quotas |
| keyenvoy | talk to keyserver |
| keyserv | server for storing public and private keys |
| mconnect | connect to SMTP mail server socker |
| named | Internet domain name server – now supported, new version |
| newkey | create a new key in the publickey database |
| nslookup | query Internet name servers |
| quotacheck | file system quota consistency checker |
| quotaon/off | turn file system quotas on and off |
| rdate | set system date from a remote host |
| repquota | summarize quotas for a file system |
| rmail | handle remote mail received via uucp |
| rquotad | remote quota server |
| tnamed | DARPA trivial name server |
| xdvstat | xdv disk statistics dynamic display |
| ypupdated | server for changing NIS information |
| zic | timezone compiler |

## Commands Changed in Appendix A

| | |
|---|---|
| cron | added information |
| dlabel | corrected example |
| exportfs | added secure NFS information/options |
| ftpd | added information/-ftp/lib |
| ifconfig | added information |
| intetd | add uid field in config file, add information |
| mailstats | print sendmail statistics |
| makedbm | added option -b for named |
| makefs | options a, d, e added. Defaults change. |
| mount | added options, secure NFS |
| ndbootd | changed information |
| nfsstat | changed information |
| portmap | changed information |
| quot | added argument |
| rdump | changed information |
| rpcinfo | added arguments |
| rrestore | changed information |
| rshd | added information |
| rstatd | added note |
| rusersd | added reference |
| setclass | corrected information re classes |
| setcomplex | corrected information re classes |
| setdefclass | corrected information re classes |
| setmaxshare | added extensive information, example |
| setsched | corrected information re classes |
| setup_client | added reference |
| showdefclass | corrected information re classes |
| srcreate | (shared regions)revised, new examples |
| srdelete | (shared regions)revised, new examples |
| srdisplay | (shared regions)revised, new examples |
| sunboot | added information |
| telnetd | 'negotiate window size' option added |
| tunefs | added -o option, new filesystem information |
| ypbind | added information |
| yppasswd | added information |
| yppush | added information |
| ypxfr | added arguments (security) |

## Procedures Added/Changed in Appendix B

| | |
|---|---|
| Add_quotas | new procedure for adding disk quotas |
| Add_user | added information on setting quotas |
| Format_disk | added information on fj2600s |
| Make_filesystem | added information on new makefs options, sizes |
| Patch_kernel | maxsharesize: corrected units to 4096 bytes.<br>usrstacksegs: default size expanded from 6 to 19 |
| Recover_System | additional information on recovering from hangs |

## RELATED DOCUMENTATION

This section lists all of the documents that you receive with your system.

| Name | Model Number |
|---|---|
| *Alliant Documentation Map* | D000.00 |
| *FX/2800 Series System Description* | D000.01 |
| *FX/2800 Series Site Preparation Guide* | D000.02* |
| *FX/800 Site Preparation Guide* | D000.03* |
| *SRM/1 Site Preparation Guide* | D000.04* |
| *FX/2800 Series Porting Guide* | D000.05 |
| *Intel i860 Programmer's Reference Guide* | D000.07 |
| *FX/2800 Series Diagnostics User's Guide* | D000.08 |
| *Concentrix-2800 Installation Guide and Release Notes* | D002.00 |
| *Concentrix-2800 Commands and Applications Manual* | D002.03 |
| *Concentrix System User's Guide* | D002.05 |
| *Concentrix Document Preparation Guide* | D002.06 |
| *GNU Emacs Manual* | D002.07 |
| *Practical Guide to UNIX* | D002.08 |
| *Learning the UNIX Operating System* | D002.09 |
| *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files and File Formats* | D002.10 |
| *Concentrix-2800 System Reference: System Calls and C Library Routines* | D002.12 |
| *Concentrix-2800 Programming Tools Guide* | D002.14 |
| *The C Programming Language* | D002.15 |
| *Scalar/C Programmer's Handbook* | D002.16 |
| *Concentrix-2800 Systems Programming Guide* | D002.17 |
| *Concentrix-2800 System Administration* | D002.18 |
| *FX/2800 Series Common Library* | D002.21 |

* you receive the appropriate manual for your site.

You can order additional copies of these manuals. You can also order these manuals in various kits. See D000.00 – *Alliant Documentation Map*.

## DOCUMENTATION CONVENTIONS

Alliant uses the following documentation conventions for the different types of text noted below:

### Command Syntax

| | |
|---|---|
| **boldface** | indicates a constant (option); type it exactly as it appears. |
| *italic* | indicates a variable; replace it with an appropriate value. |
| [ ] | (brackets) indicate an optional item or set of items. |
| { } | (braces) indicate a set of items from which you must select one. |
| ... | indicates that the previous argument-prototype can be repeated. |

### Normal Text

| | |
|---|---|
| sans serif | indicates a keyword; the name of a command, file name, or data item. |
| *italic* | indicates a variable. |
| CTRL *c* | represents a control character. You type it by holding down the control key and pressing the key for *c* at the same time. For example, to type CTRL D hold down the control key and press D at the same time. |
| RETURN | represents the return key, which is used to terminate all user input lines unless specifically indicated. |

### Examples

| | |
|---|---|
| `monospace` | smaller type indicates output from the computer. |
| `monospace` | larger type indicates input by the user. |

# CHAPTER 1
# BASIC OPERATIONS

The following tables outline the files, commands, and procedures pertaining to basic operations.

**Administration Files**

| name | description |
|---|---|
| /dev/console | console device |
| /dev/tty | terminal device * |
| /etc/fstab | static information about filesystems (edit) * |
| /etc/gettytab | terminal configuration data base (edit) * |
| /etc/group | group file (edit) * |
| /etc/mkstripe | striped partition database (edit) |
| /etc/motd | message-of-the-day file (edit) |
| /etc/nologin | file whose existence prevents logins (edit) |
| /etc/passwd | password file (edit) * |
| /etc/rc | command script for starting Concentrix-2800 |
| /etc/rc.diag | command script for starting Diagnostix-2800 |
| /etc/rc.local | site-specific command script for starting Concentrix-2800 (edit) |
| /etc/ttys | terminal initialization data (edit) * |
| /etc/utmp | login record of current users * |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /usr/adm/messages | record of messages written to console terminal |
| /usr/adm/wtmp | login record of all logins and logouts * |
| /vmunix | a symbolic link to /unix |

* See *Concentrix-2800 System Reference*: *Miscellaneous, Special Files, and File Formats*
(edit) can be edited by system administrator

**Administration Commands (Appendix A)**

| name | description |
|---|---|
| fastboot | reboot system without checking disks |
| fasthalt | halt system in console mode without checking disks |
| halt | halt system in console mode |
| init | system initialization and process control |
| rc | command scripts for starting system |
| reboot | reboot system |
| shutdown | shut down system at a given time |
| sticky | persistent text and append-only directories |

**Administration Procedures (Appendix B)**

| name | description |
|---|---|
| Boot_system | start system and bring to multi-user mode |
| Call_cs | report problem to Alliant customer service |
| Connect_diag | connect console to Diagnostix-2800 |
| Halt_system | stop system in console mode |
| Recover_system | recover from hang or panic |
| Shutdown_system | shut down system to single-user mode |

## 1.1 FRONT PANEL AND CONSOLE TERMINAL

The FX/2800 Series front panel contains a power switch (key switch), a reset button, and a power on/off light. When the key switch is in the power on (I) position: the RESET switch is enabled; and a reboot stops in console mode. When the key switch is in the power locked on (I/AUTO) position: the RESET switch is disabled; and a reboot proceeds automatically through (does not stop in) console mode. When the power is on, the light is green.

**FX/2800 Series Front Panel**

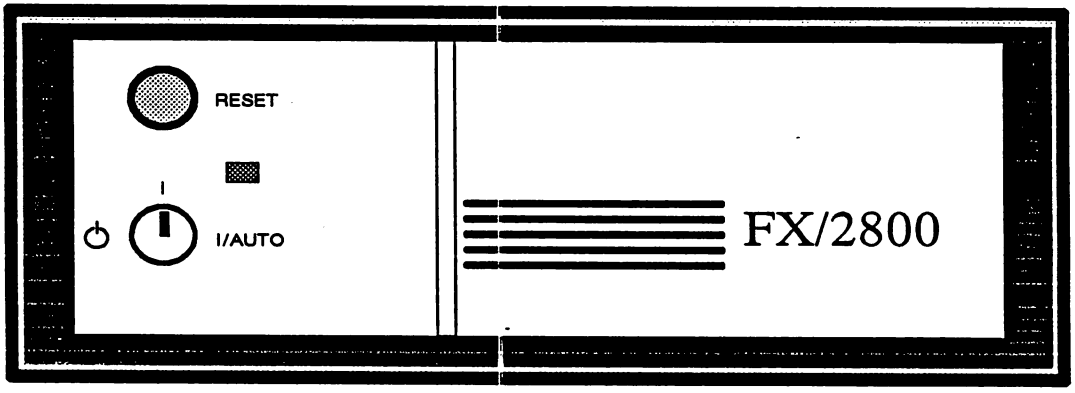


| | |
|---|---|
| RESET | Resets the hardware to a state equivalent to power-on. Clears all error registers. |
| ▒ | On/off light. |
| ⏻ | Power off position. |
| I | Power on position. |
| I/AUTO | Power locked on position. Autoboot set. Reset button disabled. |

The console terminal is a video terminal connected to the console port. The console terminal can communicate with the system in all modes of operation. Terminals not connected to the console port can only communicate with the system when the system is in Concentrix-2800 (multi-user) mode. A console printer may be connected to the console terminal to obtain hard copy of console sessions.

## 1.2 HARDWARE CONFIGURATION

The system contains processor modules and input/output modules to perform the computational work. These modules connect to one or more 512KB memory cache modules through a crossbar interconnect. Each cache module connects to a dedicated memory bus. Each memory bus connects to up to 4 memory ports (on an FX/2800 Series machine; 2 memory ports on an FX/800 Series machine); each memory port accesses half a 64MB memory module.

The following subsections provide a small amount of detail on the hardware components, so that you can recognize them in the boot script and system messages.

### 1.2.1 Processor Module

The hardware identification for a processor module (PM) is pm*n*, where *n* is the physical computational slot 0–7 (FX/2800) or 0–4 (FX/800) containing the module. Each Model 400 PM contains 4 processors, 2 concurrency control units (CCUs), and 2 cache ports. Each Model 200 PM contains 2 processors, 1 concurrency control unit (CCU), 2 256KB activated coherency caches (ACCs), and 2 cache ports.

The hardware identification for a PM processor is sce*nn*, where *nn* is a 5-bit value in hexadecimal: the 3 high-order bits identify the physical slot of the Model 400 PM or Model 200 PM; the 2 low-order bits identify the processor on the Model 400 PM (0–3) or Model 200 PM (0 or 2). The software identification for a PM processor is CE*n* (CE means computational element), where *n* is a decimal integer 0 through 1 less than the number of CEs in the machine. CE0 is the first processor on the lowest numbered PM; numbering is consecutive through the last CE on the highest numbered PM, even if the modules are not in physically adjacent slots. (The software also refers to CE*n* as ce*nn*.) On each Model 400 PM, software processor 1 is hardware processor 2, and software processor 2 is hardware processor 1; processors 0 and 3 are the same for both software and hardware. On each Model 400 PM, software processor 1 is hardware processor 2; processor 0 is the same for both software and hardware.

For example, if an FX/2800 machine has 2 Model 400 PMs, in slots 0 and 4, the configuration is as follows:

```
slot 0    slot 1    slot 2    slot 3    slot 4    slot 5    slot 6    slot 7

pm0                                     pm4
  sce00 (CE0)                             sce10 (CE4)
  sce02 (CE1)                             sce12 (CE5)
  sce01 (CE2)                             sce11 (CE6)
  sce03 (CE3)                             sce13 (CE7)
```

If an FX/2800 machine has 3 Model 200 PMs, in slots 0, 2, and 4, the configuration is as follows:

```
slot 0    slot 1    slot 2    slot 3    slot 4    slot 5    slot 6    slot 7

pm0               pm2               pm4
  sce00 (CE0)       sce08 (CE2)       sce10 (CE4)
  sce02 (CE1)       sce0a (CE3)       sce12 (CE5)
```

On each Model 400 PM, using software names, an even-numbered and an odd-numbered processor (CE*n* and CE*n*+1) connect to one CCU, and the other even-odd pair (CE*n*+2 and CE*n*+3) connect to the other CCU. On each PM, the even-numbered processors connect to one cache port, and the odd-numbered processors connect to the other cache port. The two processors that share a CCU thus access different cache ports.

On a Model 200 PM, the two processors connect to the single CCU, and each processor connects to a separate cache port.

## 1.2.2 Input/Output Module

The hardware identification for an input/output module (IOM) is iom*n*, where *n* is the physical computational slot 0–7 (FX/2800) or 0–4 (FX/800) containing the module. Each IOM contains 1 processor, 2 input/output channels, and 2 cache ports. One IOM is the master IOM and is used for starting and maintaining the system. The master IOM connects to the maintenance processor (VMEDP), which drives the diagnostics subsystem, drives the console and remote service ports, and runs Diagnostix-2800, the diagnostics operating system.

The hardware identification for an IOM processor is sce*nn*, where *nn* is a 5-bit value in hexadecimal: the 3 high-order bits identify the physical slot of the IOM; the 2 low-order bits are 00. The software identification for the processor on the master IOM is IP0 (IP means interactive processor). The software identification for a processor on a non-master IOM is IP*n*, where *n* is a decimal integer 1 through 1 less than the number of IPs in the machine. IP1 is the processor on the lowest numbered non-master IOM; numbering is consecutive through the IP on the highest numbered non-master IOM, even if the IOMs are not in physically adjacent slots. In addition, software uses machine numbers for IOMs, starting at 28 (FX/2800) or 12 (FX/800) for IP0 and decrementing by 4 for each remaining IP.

For example, if an FX/2800 machine has 3 IOMs, in slots 7 (master IOM), 3 , and 5, the configuration is as follows:

| slot 0 | slot 1 | slot 2 | slot 3 | slot 4 | slot 5 | slot 6 | slot 7 |
|---|---|---|---|---|---|---|---|
| | | | iom3 | | iom5 | | iom7 |
| | | | sce0c (IP1) | | sce14 (IP2) | | sce1c (IP0) |
| | | | (machine 24) | | (machine 20) | | (machine 28) |

## 1.2.3 Computational Configurations

One input/output module and two processor modules are included in the basic packaged FX/2800 system; one input/output module is included in the basic packaged FX/800 system. Additional input/output modules and processor modules are permitted in any combination, to a maximum of 8 (FX/2800) or 4 (FX/800) modules. Model 200 processor modules can be substituted for Model 400 processor modules. A 3-PM FX/800 system is not recommended.

The following tables list the standard configuration capabilities.

**FX/2800 (with Model 400 PMs) Computational Configurations**

| name | description |
|---|---|
| IP0 | interactive processor on master IOM – required – the master IOM connects to the VMEDP, the boot tape drive, the console, the diagnostic port, and the root disk |
| IP1 – IP7 | interactive processors on additional IOMs – optional – connects devices as desired (note: IP6 and IP7 are not typically available because the basic packaged system uses 2 of the 8 computational slots for PMs) |
| CE0 – CE7 | computational processors on 2 PMs – included in the basic packaged system |
| CE8 – CE27 | computational processors on additional PMs – optional |

**FX/2800 (with Model 200 PMs) Computational Configurations**

| name | description |
|---|---|
| IP0 | interactive processor on master IOM – required – the master IOM connects to the VMEDP, the boot tape drive, the console, the diagnostic port, and the root disk |
| IP1 – IP7 | interactive processors on additional IOMs – optional – connects devices as desired (note: IP6 and IP7 are not typically available because the basic packaged system uses 2 of the 8 computational slots for PMs) |
| CE0 – CE3 | computational processors on 2 PMs – included in the basic packaged system |
| CE8 – CE13 | computational processors on additional PMs – optional |

**FX/800 (with Model 400 PMs) Computational Configurations**

| name | description |
|---|---|
| IP0 | interactive processor on master IOM – required – the master IOM connects to the VMEDP, the boot tape drive, the console, the diagnostic port, and the root disk |
| IP1 – IP3 | interactive processors on additional IOMs – optional – connects devices as desired |
| CE0 – CE11 | computational processors on PMs – optional |

**FX/800 (with Model 200 PMs) Computational Configurations**

| name | description |
|---|---|
| IP0 | interactive processor on master IOM – required – the master IOM connects to the VMEDP, the boot tape drive, the console, the diagnostic port, and the root disk |
| IP1 – IP3 | interactive processors on additional IOMs – optional – connects devices as desired |
| CE0 – CE5 | computational processors on PMs – optional |

## 1.2.4 Memory Subsystem

The memory subsystem contains the following components:

- **Crossbar interconnect** – On the FX/2800 machine, the crossbar interconnect contains 16 data switches, 9 address switches, and 1 "mixed" switch (which actually contains the data parity bits) on 9 modules. The hardware identifications for the switches are (data) dsw0 through dswf, (address) asw0 through asw8, and (mixed) msw.

  On the FX/800 machine, the crossbar interconnect contains 8 data switches and 6 address switches on the backplane. The hardware identifications for the switches are (data) dsw0 through dsw7 and (address) asw0 through asw5.

- **Clock** – The hardware identification for the clock module is clk.

- **Cache module (CM)** – The hardware identifications for the cache modules are cm0 through cm7 (for the FX/2800; through cm1 for the FX/800).

- **Memory module (MM)** – The hardware identifications for the memory modules are (first halves) mm0a through mmfa (for the FX/2800; through mm1a for the FX/800), and (second halves) mm0b through mmfb (for the FX/2800; through mm1b for the FX/800).

Each memory module contains 2 32MB or 128MB half-modules, for a capacity of 64MB or 256MB per module and a maximum memory capacity of 4GB for an FX/2800 machine and 512MB for an FX/800 machine. The two parts of one memory module must be connected to separate cache modules.

Cache modules must be added to the system in power-of-two combinations. An FX/2800 system may have 1, 2, 4, or 8 cache modules. An FX/800 system may have 1 or 2 cache modules. Since a 1-cache system makes use of only half of each attached memory module, the use of one cache only is not recommended.

The following table lists the recommended minimum memory configurations to provide the maximum response to requests from the PMs on an FX/2800 system. On an FX/800 system, the viable options are 2 CMs (1MB or 4MB) with 1 MM (64MB or 256MB) or 2 MMs (128MB or 512MB).

**Recommended Minimum Memory Configurations**

| PMs | CMs | minimum MMs | additional MMs | increment size |
|---|---|---|---|---|
| 2 | 2 (1 or 4MB) | 1 (64 or 256MB) | up to 4 (256MB or 1GB) | 1 (64 or 256MB) |
| 3-4 | 4 (2 or 8MB) | 2 (128 or 512MB) | up to 8 (512MB or 2GB) | 2 (128 or 512MB) |
| 5-7 | 8 (4 or 16MB) | 4 (256MB or 1GB) | up to 16 (1 or 4GB) | 4 (256MB or 1GB) |

## 1.3 SOFTWARE CONFIGURATION

The system operates in various modes. The normal mode of operation is multi-user Concentrix-2800, which allows interactive use of the system by users and superusers.

### 1.3.1 Users and Superusers

System users are of two types:

- **User** – An ordinary user accesses multi-user Concentrix-2800 on a non-privileged basis. At the interactive level, a user has access to the facilities documented in the *Concentrix Commands and Applications Manual* and in Section 1 of the man (on-line documentation) database.

- **Superuser** – A superuser accesses the system on a privileged basis. At the interactive level, a superuser has access to the facilities documented in Appendix A of this manual and in Section 8 of the man database, in addition to the ordinary user facilities.

Superuser access is gained to multi-user Concentrix-2800 as follows:

- **Root account** – Log in with the user name root. The root account has a user identification of 0 and a group identification of 10 (staff). The terminal port on which you log in must be marked as secure in the file /etc/ttys (see Chapter 3, "Terminal Devices"); in this way, you can restrict superuser access to the console and other terminals, as desired, and prevent remote logins as root.

- **Superuser** – Type the command su without an argument. The user entering the su command must be a member of group wheel (group identification 0).

Such access requires knowledge of the password for the root account. A pound sign (#) appears in place of a percent sign (%) in the prompt.

In general, you should log in as root to perform administrative activities. You should log into a user account to perform non-administrative activities (prevents accidents).

### 1.3.2 File System

The file system is a hierarchical arrangement starting at a single file called the root directory (or root), proceeding through intermediate files called directories, and ending at end-node files called (plain) files.

The pathname for the root directory is a single slash (/). The pathname for other directories and plain files consists of the names separated by slashes. An absolute pathname starts with a slash (the root directory); a relative pathname does not start with a slash and denotes a pathname starting at your current position in the file system.

See Chapter 4 for system administration operations affecting the file system; see the *Concentrix System User's Guide* for a general discussion of the file system.

### 1.3.2.1 Directories Provided with Concentrix-2800

The following directories are provided with Concentrix-2800 at installation time.

**System Directories**

| name | description |
|---|---|
| / | Root directory (top of the file name tree) |
| /bin | General commands and utilities |
| /console | System files for hardware |
| /dev | Device names |
| /dist | Files typically modified by users |
| /etc | Administration commands, utilities, and other files |
| /lib | System libraries |
| /lost+found | Orphaned files located during file system check |
| /mnt | Spare mount point |
| /tmp | Temporary files |
| /usr/adm | System administration files |
| /usr/alliant | Alliant commands and utilities |
| /usr/bin | General commands and utilities |
| /usr/cs | Customer service files |
| /usr/etc | Administration commands, utilities, and other files |
| /usr/games | Games |
| /usr/hosts | Remote host names |
| /usr/include | Parameter files for user programs |
| /usr/lib | User libraries |
| /usr/local | Site-specific commands and utilities |
| /usr/man | Man pages (on-line documentation) |
| /usr/msg | System-wide message facility |
| /usr/pub | Usage unknown |
| /usr/skel | Default login scripts |
| /usr/spool | Storage for printer, mail, and UUCP data |
| /usr/src | Files for regenerating system software |
| /usr/sys5 | System V utilities |
| /usr/tmp | Temporary files |
| /usr/ucb | Extensions to commands and utilities |

### 1.3.2.2 Files that Grow

The following system files grow with system use and their disk usage should be monitored. As necessary, you should archive the contents of these files and reduce their sizes (to zero bytes, if desired). Do not, however, remove these files entirely.

**System Files that Grow**

| name | description |
|---|---|
| /etc/utmp | Record of current system users |
| /usr/adm/acct | Execution accounting file |
| /usr/adm/aculog | Error log for tip and cu |
| /usr/adm/lastlog | Login activity for last command |
| /usr/adm/lpd-errs | Line printer daemon errors |
| /usr/adm/savacct | Summary of execution accounting file |
| /usr/adm/usracct | Per-user summary of execution accounting file |
| /usr/adm/messages | Record of messages written to console terminal |
| /usr/adm/wtmp | Record of all logins and logouts |
| /usr/lib/batch/log | Batch system log |
| /usr/spool/uucp/ERRLOG | UUCP error log |
| /usr/spool/uucp/LOGFILE | UUCP log file |
| /usr/spool/uucp/SYSLOG | UUCP system log |

The following example moves /usr/adm/messages to another file system that serves as an archive, then creates a new empty messages file:

```
# cd /usr/adm
# mv messages /archives/usr/adm/messages.7_13_90
# touch messages
```

The next example simply deletes all but the last 250 lines of /usr/adm/messages:

```
# cd /usr/adm
# tail -250 messages > tmp
# mv tmp messages
```

### 1.3.2.3 Files that Need Maintenance

The following system files need periodic attention depending on the activities at your site.

**System Files That Need Periodic Maintenance**

| name | description |
|---|---|
| /console/console.config | configuration of your machine |
| /etc/disktab | disk description file |
| /etc/ethers | ethernet address data base |
| /etc/fstab | static information about file systems |
| /etc/gettytab | terminal configuration data base |
| /etc/group | group file |
| /etc/hosts | host name data base |
| /etc/hosts.equiv | list of trusted hosts |
| /etc/inetd.conf | configuration file for the inetd daemon |
| /etc/motd | message of the day |
| /etc/networks | network name data base |
| /etc/passwd | password file |
| /etc/printcap | printer capability data base |
| /etc/rc.local | site-specific command script for starting Concentrix-2800 |
| /etc/qconf | batch configuration file |
| /etc/remote | remote host description file |
| /etc/services | service name data base |
| /etc/syslog.conf | configuration file for the syslogd daemon |
| /etc/ttys | terminal initialization data |
| /usr/lib/aliases | aliases file for sendmail |
| /usr/lib/crontab | commands run periodically by cron |
| /usr/lib/sendmail.cf | sendmail configuration file |

### 1.3.2.4 Listing Files and Directories

A useful command for listing the contents of a directory is the ls command with the –lt option. This command lists the files in the directory in long format starting with the most recently modified file. Long format provides the following information (reading right to left): name of the file, last time of modification, size in bytes, name of owner, number of links, and permissions.

Note the following example:

```
% ls -lt /usr/adm
total 3686
-rw-rw-r--  1 root       132320 Mar 22 16:18 acct
-rw-rw-rw-  1 root       461052 Mar 22 16:17 wtmp
-rw-rw-rw-  1 root        65838 Mar 22 16:13 lpacct
-rw-rw-rw-  1 root        28112 Mar 22 16:13 lastlog
-rw-rw-r--  1 root         8656 Mar 22 15:49 maillog
-rw-r--r--  1 root        39645 Mar 22 15:49 messages
-r--------  1 root          311 Mar 22 03:30 check.today
-rw-rw-rw-  1 root        28364 Mar 22 03:26 usracct
-rw-rw-r--  1 root        60534 Mar 22 03:26 savacct
-rw-r--r--  1 root       299728 Mar 17 18:07 messages.0
-r--------  1 root          311 Mar 16 04:06 check.yesterday
-rw-rw-r--  1 root          131 Mar 13 11:13 aculog
-rw-r--r--  1 root       281619 Mar 11 02:22 messages.1
-rw-rw-rw-  1 root          846 Mar  6 09:51 lpd-errs
-rw-r--r--  1 root       306012 Mar  3 18:09 messages.2
-rwxr-xr-x  1 root         3726 Mar  2 18:43 daily
-rw-rw-r--  1 root      1919988 Mar  1 05:28 wtmp.0
-rw-rw-r--  1 root         7224 Feb 23 11:39 routelog
-rwxr-xr-x  1 root          834 Feb 21 11:05 weekly
-rwxr-xr-x  1 root          530 Feb 21 11:05 monthly
-rw-rw-r--  1 root            0 Feb 21 07:21 lptest-errs
```

## 1.3.3 Modes of Operation

The system operates in various modes, as outlined in the following table and in the following subsections. The system startup procedure passes through a mode automatically unless you set a flag to stop in that mode. Flags are set by the options command in console mode; and by the halt (fasthalt), reboot (fastboot), and shutdown commands. You can also set the machine to power on (I), rather than power locked on (I/AUTO), to stop in console mode.

**Modes of Operation**

| flag | mode | how to stop in | prompt | how to continue |
|---|---|---|---|---|
| H | console | power on RESET<br>halt -H<br>reboot -H<br>shutdown -H | cons> | options = [bsdBSD]RETURN unix |
| B | VMEbus booter | console option B<br>reboot -B<br>shutdown -B | boot from: dunix:<br>boot from: dddt:<br>root device: | [*alternate dunix*]RETURN<br>[*alternate dddt*]RETURN<br>[*alternate root device*]RETURN |
| D | Alliant ddt | console option D<br>reboot -D<br>shutdown -D | none | ESCAPE p or ESCAPE g |
| S | Diagnostix-2800 | console option S<br>halt*<br>reboot -h*<br>reboot -S<br>shutdown -h*<br>shutdown -S<br>*Shuts down to Diagnostix-2800; does not reboot. | dunix-sh# | CTRL d |
| | FX/2800 startup | as above | Reboot with ...: | y RETURN (reboots)<br>t RETURN (connects Concentrix)<br>h RETURN (connects Diagnostix) |
| b | FX/2800 booter | console option b<br>reboot -b<br>shutdown -b | unix:<br>ddt:<br>cfg: | [*alternate unix*]RETURN<br>[*alternate ddt*]RETURN<br>[*alternate console.config*]RETURN |
| d | Alliant ddt 2800 | console option d<br>reboot -d<br>shutdown -d | none | ESCAPE p or ESCAPE g |
| s | single-user | console option s<br>reboot -s<br>shutdown*<br>shutdown -s<br>*Shuts system down to single-user; does not reboot. | Concentrix-sh# | CTRL d<br>CTRL ^! (escapes to Diagnostix-2800)<br>CTRL ^. (escapes to FX/2800 startup) |
| | Concentrix-2800 | reboot<br>shutdown -r | *system(n)*# | log into Concentrix-2800<br>CTRL ^! (escapes to Diagnostix-2800)<br>CTRL ^. (escapes to FX/2800 startup) |

Upon completion of startup, two operating systems are running: single-user Diagnostix-2800 is running on the VMEDP in FX/2800 startup mode; multi-user Concentrix-2800 is running on the IPs and CEs. Because the VMEDP and the main processors are of different types, each operating system (Diagnostix-2800 and Concentrix-2800) has its own root file system on its own disk partition and cannot access each other's files.

The console is connected to Concentrix-2800 from Diagnostix-2800 through the tip (connect to a remote system) utility. At this point, you can log into Concentrix-2800, for example, as root, on the console. You can connect the console to Diagnostix-2800 (and back to Concentrix-2800) as explained below under "FX/2800 Startup Mode."

### 1.3.3.1 Console Mode

The console program executes from an EPROM (erasable programmable read-only memory) on the VMEDP and initiates execution of the VMEbus booter. Console mode is entered as follows:

- **Power on** – The machine is turned from the power off to the power on or power locked on (I/AUTO) position. The power on position causes the machine to stop in console mode. The power locked on (I/AUTO) position causes the machine to proceed through console mode (reboot automatically).

- **Reset** – The RESET switch is pressed while the machine is in the power on position. Keeping the machine in the power on position causes the machine to stop in console mode. Turning the machine to the power locked on (I/AUTO) position causes the machine to proceed as if u were typed in response to the cons> prompt.

- **Software reboot** – The reboot or shutdown command is typed from Concentrix-2800 or single-user mode with the -B, -D, or -S option; or the reboot command is typed from Diagnostix-2800 with any option except -H. The machine stops in console mode if the machine is in the power on position. The machine proceeds through console mode if the machine is in the power locked on (I/AUTO) position.

- **Halt** – The halt, reboot, or shutdown command is typed from Concentrix-2800 or single-user mode with the -H option. The halt or reboot command is typed from Diagnostix-2800 with the -H option. The machine stops in console mode.

If the machine stops in console mode, the following prompt appears:

```
cons>
```

The most common console commands are as follows:

- **options** – The options command displays and sets the boot flags, the flags that stop processing in other modes during system startup. Type o; the system fills out the word options, displays the current flags, and explains basic choices. A blank display indicates that no flags are set. Type any combination of BDSbds, followed by a carriage return, to set flags for the specified modes (type an invalid option to delete current flags; respond with ? to obtain a description of the flags). The following command sets the boot flags to stop in single-user mode, then resets them to stop in both Diagnostix-2800 and single-user modes:

```
cons>options = ; Enter RETURN to keep, SPACE to clear, or new options.
New Options? s[RETURN]

cons>Options = s Enter RETURN to keep, SPACE to clear, or new options.
New Options? Ss[RETURN]

cons>
```

- **unix** – The unix command causes system startup to commence. Type u; the system fills out the word unix and processing continues. From a remote terminal, you must set parameters 6[RETURN] before typing unix.

```
cons>unix
```

- **parameters** – The parameters command enables and disables the remote port. Type p; the system fills out the word parameters and a menu appears. Type 6 to enable the remote port and 7 to disable it.

At any time, type a question mark in response to the cons> prompt to obtain a list of console commands. The following example demonstrates the complete list of console commands:

```
cons>?

Commands are:
        B/W/L   - byte/word/long mode for examinations
        C       - enable IPC clock
        E/S     - open address in supervisor/selected function code space
        F       - set boot flags
        G       - transfer execution to address
        I/M     - initialize map to global/local memory
        P       - set parameters (console/remote baud rates, boot devices)
        R       - reset system
        T       - run tests
        U       - boot
        Dbcu    - device boot from specific unit
        ?       - this list

An open location may be modified by entering a new value
An open location is terminated by: <cr> - close, <lf> - next, ^ - previous
A delete cancels the current command
An address ending with \ opens without read
All numbers are in hex
cons>
```

### 1.3.3.2 VMEbus Booter Mode

The VMEbus booter loads the Diagnostix-2800 kernel (by default, the executable file /dunix on xdv00p) and the Diagnostix-2800 kernel debugger (by default, the executable file /dddt on xdv00p), and initiates execution of Diagnostix-2800 on the VMEDP.

VMEbus booter mode follows console mode in the startup process. Processing stops in VMEbus booter mode if the B flag is set from console mode, or reboot -B or shutdown -B is in effect. Otherwise, processing proceeds.

If you stop in VMEbus booter mode, prompts appear for an alternate kernel, debugger, and root device (/dev/xdv00p by default). Respond as follows:

RETURN — to specify the default. The defaults are the files dunix and dddt, on the first partition labeled as the Diagnostix-2800 root partition (dlabel command, -Dr*p* for Diagnostix-2800 root device). The search order for the first root partition is disk controllers and units in ascending order in the first VME chassis.

xdv(*c*,*u*,*p*)*filename* — to specify an alternate kernel image or debugger, where *c* is the controller number; *u* is the unit number; and *p* is the partition letter. For example: xdv(0,0,p)dunix or xdv(0,0,p)dddt. The disk specification can be omitted if the default is acceptable, for example, myunix to specify another kernel on the default device. The file name can be omitted if the default is acceptable, for example, xdv(1,0,p) to specify a kernel with the default file name on another device.

Omit *filename* in response to the root device prompt.

%h — to display help text.

%ls — to list the files in the root partition of the current device.

%r — (applies only to dunix prompt) to read a ramdisk from a device.

ESCAPE — (applies only to dddt prompt) to skip loading a Diagnostix-2800 debugger.

The following example specifies an alternate kernel, and uses the defaults for the debugger and the root device:

```
Alliant VMEbus Booter [made: Wed May 9 19:32:23 EDT 1990]
Type %h for help

boot from:  xdv(0,0,p)dunix: mydunixRETURN
loaded @ 0x0: 469048.+63272.+50864.+38688.+39264. start @ 0x0
boot from:  xdv(0,0,p)dddt: RETURN
map_mem: start=670000, ddt_addr=3d0000, size=31, flags=a0000001
loaded @ 0x670000: 78856.+7064.+37496. start @ 0x670000
root device:  xdv(0,0,p): RETURN
```

Processing continues after your final response.

### 1.3.3.3 Alliant ddt Mode

Diagnostix-2800 executes under control of the kernel debugger ddt unless you suppress the debugger in VMEbus booter mode (typing ESCAPE in response to the dddt prompt).

Alliant ddt mode follows VMEbus booter mode in the startup process. Processing stops in Alliant ddt mode if the D flag is set from console mode, or reboot -D or shutdown -D is in effect. Otherwise, processing proceeds.

The debugger does not prompt if you stop in Alliant ddt mode. A header line starting with the words Alliant ddt appears and processing stops. At this point, you can enter ddt commands; see the *Concentrix-2800 Systems Programming Guide*. To continue processing, type ESCAPE g or ESCAPE p (note: ddt echoes ESCAPE as a dollar sign). In the following example, the user stops in Alliant ddt mode and immediately continues:

```
Alliant ddt X.82 M cpu_id=0. [lnk: Wed Mar  1 16:03:53 EST 1989]
ESCAPE p
```

### 1.3.3.4 Diagnostix-2800 Mode

Diagnostix-2800 is an operating system that runs on the VMEDP and provides access to the diagnostics subsystem. This operating system remains running during the time that Concentrix-2800, the main operating system, is running. At startup time, Diagnostix-2800 executes the script /etc/rc.diag. This script checks and mounts the root file system, then enters FX/2800 startup mode by invoking the script /console/console_loop.

Diagnostix-2800 mode follows Alliant ddt mode in the startup process. Processing stops in Diagnostix-2800 mode if the S flag is set from console mode, or reboot -S or shutdown -S is in effect. Otherwise, processing proceeds.

You can shut the system down to Diagnostix-2800 mode with one of the following commands: halt, reboot -h, or shutdown -h. A shutdown to Diagnostix-2800 terminates Concentrix-2800 and connects the console to Diagnostix-2800, which remains running.

You can connect to (obtain a shell in) Diagnostix-2800 mode through the h option of FX/2800 startup mode (see below) and the ! (exclamation mark) command of the FX/2800 booter (see below).

After Concentrix-2800 is activated, you can connect to (obtain a shell in) Diagnostix-2800 mode from Alliant ddt 2800 mode, single-user mode, or Concentrix-2800 (multi-user) mode by typing CTRL ^ (hold down the control and shift keys, then press the 6 key) followed by an exclamation mark, as demonstrated below:

```
Concentrix(118)# CTRL ^ !
dunix(1)#
```

If you stop in Diagnostix-2800 mode during startup, or shut down or connect to Diagnostix-2800 mode, the prompt dunix(1)# or dunix-sh# appears. At this point, you can enter commands to the Diagnostix-2800 operating system. Type CTRL d to continue startup or reconnect to the previous mode, as shown below:

```
dunix(9)# CTRL d
```

### 1.3.3.5 FX/2800 Startup Mode

At startup time, the Diagnostix-2800 script /console/console_loop (which initiates FX/2800 startup mode) executes the diagnostic monitor on the FX/2800 hardware, executes the FX/2800 booter, starts Concentrix-2800 under debugger control, and connects the console to Concentrix-2800 through the tip (connect to a remote machine) command.

FX/2800 startup mode follows Diagnostix-2800 mode in the startup process.

If processing stops in Diagnostix-2800 mode during startup (S flag or -S option), processing also stops in FX/2800 startup mode; otherwise, processing proceeds. After you type CTRL d to exit Diagnostix-2800 (if you stop in this mode), you are prompted on how to proceed:

```
FX/2800 Startup

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

Respond y RETURN to continue system startup.

After Concentrix-2800 is activated, you can disconnect from Alliant ddt 2800 mode, single-user mode, or Concentrix-2800 (multi-user) mode by typing CTRL ^ (hold down the control and shift keys, then press the 6 key) followed by a period. This action connects you to FX/2800 startup mode (without affecting the Concentrix-2800 mode), as demonstrated below:

```
Concentrix(118)# CTRL ^.
[EOT]

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

To obtain a shell in Diagnostix-2800 mode, respond to the prompt with a lowercase h followed by a carriage return:

```
Reboot with flags set to "b"? (y for yes, ? for help): h RETURN
dunix(1)#
```

To terminate the Diagnostix-2800 shell and return to FX/2800 startup mode, type CTRL d in response to the dunix prompt:

```
dunix-sh# CTRL d
Reboot with flags set to ""? (y for yes, ? for help):
```

To reconnect to the previous mode, respond to the prompt with a lowercase t followed by a carriage return. Type another carriage return after you are connected.

```
Reboot with flags set to ""? (y for yes, ? for help): t[RETURN]
restarting tip
connected
[RETURN]
Concentrix(118)#
```

At any time in FX/2800 startup mode, you can respond to the prompt with a question mark to display the full list of options:

```
Reboot with flags set to ""? (y for yes, ? for help): ?

Response can be a combination of these characters:

 s - reboot concentrix, stop in single user concentrix
 b - reboot concentrix, stop in the concentrix booter
 d - reboot concentrix, stop in concentrix ddt
 h - "halt" (a dunix shell which, on exit, returns to the "Flags" prompt)

Or any combination of these characters:

 S - reboot dunix, stop in single user dunix
 B - reboot dunix, stop in the dunix booter
 D - reboot dunix, stop in dunix ddt
 H - halt dunix in IP cons> prompt

Or one of these single characters

' '  (space) reboot concentrix with no flags
 t - don't reboot, just restart tip
 y - reboot concentrix with shown flags
 R - enable / disable booting from ramdisk
 X - exit console_loop, go to single user Dunix (via 'kill -TERM 1')

Reboot with flags set to ""? (y for yes, ? for help):
```

### 1.3.3.6 FX/2800 Booter Mode

The FX/2800 booter loads the Concentrix-2800 kernel (by default, the executable file /unix on xdv00r) and the Concentrix-2800 kernel debugger (by default, the executable file /ddt on xdv00r), and initiates execution of Concentrix-2800 in single-user mode on IP0.

FX/2800 booter mode follows FX/2800 startup mode in the startup process. Processing stops in FX/2800 booter mode if the b flag is set from console mode, or reboot -b or shutdown -b is in effect. Otherwise, processing proceeds.

If you stop in FX/2800 booter mode, prompts appear for an alternate kernel, debugger, and console configuration file (/etc/console.config in the Concentrix-2800 root partition or /console/console.config in the Diagnostix-2800 root partition, by default). Respond as follows:

RETURN — to specify the default. The defaults are the files unix and ddt, on the first partition labeled as the Concentrix-2800 root partition (dlabel command, -C*rp* for Concentrix-2800 root device). The search order for the first root partition is disk controllers and units in ascending order in the first VME chassis.

/dev/rxdv*cup*:/*filename* — to specify an alternate kernel image or debugger, where *c* is the controller number; *u* is the unit number; and *p* is the partition letter. For example: /dev/rxdv00r:/unix or /dev/rxdv00r:/ddt. Full names must be specified. Specify only *filename* in response to the cfg prompt.

! — to suspend the FX/2800 booter and obtain a shell in Diagnostix-2800 mode. Type CTRL d to terminate the Diagnostix-2800 shell and return to the FX/2800 booter.

The following example specifies an alternate kernel, and uses the defaults for the debugger and the console configuration file:

```
FX/2800 Booter

unix [/dev/rxdv00r:/unix]: /dev/rxdv00r:/myunix RETURN
ddt [/dev/rxdv00r:/ddt]: RETURN
cfg [console.config]: RETURN
```

Processing continues after your final response.

### 1.3.3.7 Alliant ddt 2800 Mode

Concentrix-2800 executes under control of the kernel debugger ddt.

Alliant ddt 2800 mode follows FX/2800 booter mode in the startup process. Processing stops in Alliant ddt 2800 mode if the d flag is set from console mode, or reboot -d or shutdown -d is in effect. Otherwise, processing proceeds.

The debugger does not prompt if you stop in Alliant ddt 2800 mode. A header line starting with the words Alliant ddt 2800 appears and processing stops. At this point, you can enter ddt commands; see the *Concentrix-2800 Systems Programming Guide*. To continue processing, type ESCAPE g or ESCAPE p (note: ddt echoes ESCAPE as a dollar sign). In the following example, the user stops in Alliant ddt 2800 mode and immediately continues:

```
Alliant ddt 2800 X.82 M cpu_id=0. [lnk: Wed Mar  1 16:03:53 EST 1989]
ESCAPE p
```

You can disconnect from Alliant ddt 2800 mode and connect to Diagnostix-2800 mode by typing CTRL ^! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from Alliant ddt 2800 mode and connect to FX/2800 startup mode by typing CTRL ^. (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

### 1.3.3.8 Single-User Mode

In single-user mode, Concentrix-2800 executes a subset of the full services and permits only superuser access from the console.

Single-user mode follows Alliant ddt 2800 mode in the startup process. Processing stops (you are given a shell) in single-user mode if the s flag is set from console mode, or reboot -s or shutdown -s is in effect. Otherwise, processing proceeds.

You can shut the system down to single-user mode with the following command: shutdown.

If you stop in single-user mode during startup, or shut down to single-user mode, the prompt Concentrix-sh# appears. At this point, you can enter commands to the Concentrix-2800 operating system. Type CTRL d to continue processing (that is, to start multi-user Concentrix-2800), as shown below:

```
Concentrix-sh# CTRL d
```

You can disconnect from single-user mode and connect to Diagnostix-2800 mode by typing CTRL ^! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from single-user mode and connect to FX/2800 startup mode by typing CTRL ^. (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

### 1.3.3.9 Concentrix-2800 Mode

In Concentrix-2800 (multi-user) mode, Concentrix-2800 executes all services and permits users to log in on terminals, pseudo-terminals, and the console.

Concentrix-2800 mode follows single-user mode in the startup process. The system is fully started when the login prompt appears on the console:

```
ALLIANT CONCENTRIX (mysystem)

login:
```

At this point, you can use the console as a login terminal.

You can disconnect from Concentrix-2800 mode and connect to Diagnostix-2800 mode by typing CTRL^! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from Concentrix-2800 mode and connect to FX/2800 startup mode by typing CTRL^. (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

## 1.4 START-UP AND SHUTDOWN

The procedures for starting and restarting the system differ depending on the current state of the system and your planned activities:

- **Reboot from multi-user mode** – Use the shutdown command with the –r option, as described in Procedure A of Boot_system (Appendix B).
- **Halt from multi-user mode** – Use the shutdown command with the –H option, as described in Halt_system (Appendix B). You should always halt the system before shutting off the power.
- **Shut down to single-user mode** – Use the shutdown command as described in Shutdown_system (Appendix B). From single-user mode, you can reboot the system, halt the system, or exit to multi-user mode as described in Procedure B of Boot_system (Appendix B). You can also shut down by sending a terminate signal to the init process (see init in Appendix A), but this procedure should only be used in emergency and test situations.
- **Shut down to Diagnostix-2800 mode** – Use the shutdown command with the –h option. This action shuts down Concentrix-2800 and connects the console to Diagnostix-2800 on the VMEDP. From Diagnostix-2800 mode, you can reboot the system, halt the system, or exit to FX/2800 startup mode.
- **Power up** – Use the power switch on the front panel as described in Procedure D of Boot_system (Appendix B). In general, you should not power up unless the machine is powered down (for some reason such as hardware maintenance) or the machine is hung and does not restart properly from reset.
- **Reset** – Use the reset switch on the front panel as described in Procedure C of Boot_system (Appendix B). In general, you should not reset the machine unless the entire system (including Diagnostix-2800 is hung).

- **Reboot from Diagnostix-2800** – Connect to FX/2800 startup mode and select a reboot option as described in Procedure E of Boot_system (Appendix B). In general, you should not reboot from Diagnostix-2800 unless Concentrix-2800 is hung.

## 1.4.1 Shutdown Versus Reboot

The reboot and fastboot commands take the system away from interactive users without warning and should not be used in multi-user mode. Use shutdown –r or shut down to single-user mode and then issue the reboot or fastboot command.

## 1.4.2 Fastboot

The file system is checked for consistency as part of normal start-up; see Chapter 4 for detailed information. The file checks can take awhile depending on the amount of secondary storage on the system.

The command fastboot, shutdown –f, or the existence of a file named /fastboot when a reboot occurs, bypasses the file system checks. Bypassing the file system checks to save time is safe if the system is operating correctly and is shut down in an orderly manner. The recommended procedure is to shut down to single-user mode, then issue the fastboot command.

Never bypass the file system checks when rebooting after a system failure or if any problem is suspected in the file system.

## 1.4.3 Boot Device

To start the system: the console program (console mode) must load the booter from disk or tape; the booter (VMEbus booter mode) must load Diagnostix-2800 from the Diagnostix-2800 root partition; and Diagnostix-2800 (FX/2800 booter mode) must load Concentrix-2800 from the Concentrix-2800 root partition. Default search rules exist for locating the device(s) containing these programs, and you can respecify the boot device(s). Typically, these programs exist on Xylogics disk controller 0, unit 0, or – if the system has no Xylogics disk – SCSI disk bus 0, controller 0.

### 1.4.3.1 Booter Program

In console mode, specify the starting point for searching for the boot device using one of the following commands:

- u command to start the search at the current default established by previous D command.
- d command to start the search with a given type of device without changing the default for subsequent booting.
- D command to make your specification the default for the u command for subsequent booting.

To use the d or D command, (d in this example) specify dd*cu* for Xylogics disks, ds*bu* for SCSI disks, and dt*bu* for SCSI tapes. The console program fills in and prompts as you type the letters, and tells you which controller and unit it is trying, for example:

```
cons>device boot from
(d)isk/(t)ape/(s)scsi = d
controller (0-7) [return for default] = 0
unit (0-7) = 0
xdv0**: 754, prom=2282, firmware=2.4.0
------ Trying disk controller(0) unit(0).
   xdv00*: "fj690s" <624,27,64+4> [512 b/s]
Booter @100000: text=20AC8 + data=7550 + bss=1C278
```

The search for the boot device begins with the specified unit, then proceeds to the other units of the same device type in ascending order. If you specify dd or ds, the search proceeds (if necessary) to the other disk type, then to the tape devices. If you specify dt, the search proceeds (if necessary) to the Xylogics disks, then to the SCSI disks.

If no D command has been entered to specify the desired default (unlikely), the console program (console mode) starts and continues the search for the booter program in the following order when you enter the u command:

1. The Xylogics disk (xdv) on controller 0, unit 0.
2. The remaining Xylogics disks attached to the master IOM (IP0) in ascending order by controller and unit.
3. The SCSI disk (sd) on bus 0, SCSI id 0.
4. The remaining SCSI disks attached to the master IOM (IP0) in ascending order by bus and unit.
5. The SCSI tape (sch) on bus 0, SCSI id 2.

### 1.4.3.2 Diagnostix-2800 Root Partition

Following the search order explained above, the partition marked as Diagnostix-2800 root (typically p) found on the same disk the booter was read from is selected and announced as the "assigned root." If no marked root partition is found on that disk, the first partition encountered with a "Diagnostix-2800 root" marking is chosen and announced as "tentative" root, and any subsequent partitions so marked are announced as "possible" roots. The user can respecify the Diagnostix-2800 root partition by stopping in VMEbus booter mode.

The following example is a continuation of the printout from the d command entered above:

```
Alliant VMEbus Booter [made: Sun Dec 9 20:16:04 EST 1990]
Configuring devices ...
```

*Note the new lines "assigned root" and "possible root."*

```
xdv0**: 754, prom=2282, firmware=2.4.0
   xdv00*: "fj690s" <624,27,64+4> [bsize=512]
   assigned  root (from console proms): xdv00p
vme_scsi@0x6fc000: code(091,0) firmware(XAN,06/26/90)
   sd00*: "IMP94601", 0 MB, 1024 b/s
   exabyte[0,2]: EXABYTE EXB-8200        4.24 sequential (removable)
vme_scsi@0x6fb800: code(077,0) firmware(XAL,02/20/90)
   sd24*: "SONSMO-C", 311 MB, 1024 b/s
   possible  root: sd24p
   sd26*: " SE      ", 20 MB, 512 b/s
   possible  root: sd26h

Type %h for help

boot from:  xdv(0,0,p)dunix:
```

### 1.4.3.3 Concentrix-2800 Root Partition

The Concentrtix-2800 booter uses the disk containing the current Diagnostix-2800 root partition as the starting point to search for a partition marked "Concentrix-2800 root" (typically r). If none is found on the first disk, it proceeds to the other units and controllers of the same device type in ascending order, wrapping to the device at 0,0, and returning to the starting point. It then proceeds to the "other" disks (xdv or sd) and searches starting with 0,0. The user can respecify the Concentrix-2800 root partition by stopping in FX/2800 booter mode.

## 1.4.4 Boot Script

As start-up proceeds, a large number of messages appear on the console describing events; these messages are called the boot script. Your console printer should be turned on to log the boot script. A large portion of the boot script is also logged in the file /usr/adm/messages.

Below is an example of a boot script. The boot script that you see on your system will vary from this example because the processor configuration and the exact list of devices attached to the system will probably be different.

Depending on how the boot flags are set, the procedure may stop at various points for input. Interspersed in the script, in italics, are remarks that explain the events and note where responses may be required; see "Modes of Operation" above and Boot_system in Appendix B for more detailed information. The example starts after pushing the reset button.

*If starting from power on, hardware self-tests are run.*

```
Reset...Alliant Computer Systems Corporation
Diagnostic Processor (Code Rev 1.35  4-Mar-90 17:20)
Enabling map
```

*The system stops in console mode if the H flag is in effect or if the power switch is turned to 'power on' rather than 'power locked.' To see the options available, type 'o' at the prompt, then '?'*

```
cons>Options = r; Enter RETURN to keep, SPACE to clear, or new options.
New Options? ?

The options field is an alternate (symbolic) form of the flags
The options are:
        B - stop in the Dunix booter and ask for file names for boot
        D - stop in Dunix ddt and wait for commands
        S - stop in single user mode Dunix

        b - stop in the Concentrix booter and ask for file names for boot
        d - stop in Concentrix ddt and wait for commands
        s - stop in single user mode Concentrix

        r - remote diagnostics port is enabled
cons>
```

*If you type 'o' then BDSbds[RETURN], the booter stops at all the flags.*

```
cons>Options = ; Enter RETURN to keep, SPACE to clear, or new options.
New Options? BDSbds[RETURN]
```

*Type 'u' to start the system.*

```
cons>unix
xdv0**: 754, prom=2282, firmware=2.4.0
------ Trying disk controller(0) unit(0).
   xdv00*: "fj690s" <624,27,64+4> [512 b/s]
Booter @100000: text=21408 + data=76B8 + bss=1C288

Alliant VMEbus Booter [made: Wed Dec 19 19:32:23 EDT 1990]
Configuring devices ...
xdv0**: 754, prom=2282, firmware=2.4.0
   xdv00*: "fj690s" <624,27,64+4> [bsize=512]
   assigned  root (from console proms): xdv00p
vme_scsi@0x6fc000: code(091,0) firmware(XAN,06/26/90)
   sd00*: "IMP94601", 1035 MB, 1024 b/s
   exabyte[0,2]: EXABYTE EXB-8200        4.24 sequential (removable)
vme_scsi@0x6fb800: code(077,0) firmware(XAL,02/20/90)
   sd24*: "SONSMO-C", 311 MB, 1024 b/s
   possible  root: sd24p
   sd26*: " SE      ", 20 MB, 512 b/s
   possible  root: sd26h
Type %h for help
```

*The system stops in VMEbus booter mode if the B flag is in effect, and prompts for a kernel name, a debugger name, and the root directory for Diagnostix-2800. Type [RETURN] to accept the defaults. Type the names of the replacements if you want to use other ones. Press [ESCAPE] in response to the debugger prompt if you do not want a debugger. You can type %h for help or %ls for a listing of the files in the root directory.*

```
boot from:  xdv(0,0,p)dunix: [RETURN]
loaded @ 0x0: 469128.+63272.+50832.+38640.+39232. start @ 0x0
boot from:  xdv(0,0,p)dddt: [RETURN]
map_mem: start=670000, ddt_addr=3d0000, size=31, flags=a0000001
loaded @ 0x670000: 78856.+7064.+37496. start @ 0x670000
root device:  xdv(0,0,p): [RETURN]
Alliant ddt X.82 M cpu_id=0. [lnk: Wed Mar  1 16:03:53 EST 1989]
```

*The system stops in Alliant ddt mode if the D flag is in effect; the debugger does not prompt. Type* ESCAPE *p or* ESCAPE *g to proceed; type :e to reboot.*

```
ESCAPE g
ip00>

*-------------------------------------------------------------------------*
|                      Alliant Diagnostix  4.B2d                          |
*-------------------------------------------------------------------------*
|                     Thu Jun 7 14:21:29 EDT 1990                         |
| Copyright (c) 1985 Alliant Computer Systems Corporation (Unpublished)   |
|                        All Rights Reserved                              |
*-------------------------------------------------------------------------*
dp00> Real Memory: 976 pages, 3.8 Mb
dp00> Palloc Memory: 350 pages, 1.4 Mb
dp00> Available Memory: 666 pages, 2.6 Mb
dp00> Using 97 buffers containing: 97 pages, 0.4 Mb
dp00>
Configuring VME IP0 (0x1801)
dp00> ex[0]: online, versions: nx 5.5, ex 2.0
dp00> ex[0]: ex[0]: interrupt vector = 0x40
dp00> xdv0**: 754, prom=2282, firmware=2.4.0
dp00> install: root device (from booter) = xdv00p
dp00> install: swap: xdv00q
dp00> xdv00*: "fj690s" <624.,27.,64.> spares=4
dp00> xdv1**: 754, prom=2282, firmware=2.4.0
dp00> xdv10*: "fj1000s" <745.,27.,80.> spares=2
dp00> vme_scsi@0x6fc000: code(077,0) firmware(XAL,02/20/90)
dp00> exabyte[2,0]: EXABYTE EXB-8200        4.24 sequential (removable)
dp00> mb_init: cpu 0. max iomaps=955. 3816Kb transfers possible.
dp00> soft_net_cpu, 0, doesn't exist.  Changing to 0.
dp00> Paging File space: 1425 pages, 5.6 Mb
dp00> Swapping File space: 1166 pages, 4.6 Mb
dp00> Available user memory: 543 pages, 2.1 Mb
```

*The system stops in Diagnostix-2800 mode if the S flag is in effect. Type* CTRL d *to continue.).*

```
dunix-sh# CTRL d
```

*The system stops in FX/2800 startup mode when you exit from Diagnostix-2800 mode (if you stop in that mode).*

```
dunix-sh# Thu Jul 12 17:49:15 EST 1990
FX/2800 Startup

Entering Console Boot Loop.

Reboot with flags set to "bds"? (y for yes, ? for help): y RETURN
```

*The system stops in the FX/2800 booter mode if the b flag is in effect, and prompts for a kernel name, a debugger name, and the console.config file for Concentrix-2800. Type* RETURN *to accept the defaults. Type the names of the replacements if you want to use other ones. You can type ! to return to Diagnostix-2800 mode.*

```
FX/2800 Booter - Selecting boot files
(At each prompt, ! gives sh, ^C<cr> returns to flags prompt)
unix [unix]: RETURN
ddt [ddt]: RETURN
cfg [console.config]: RETURN
```

*The diagnostics monitor initializes the FX/2800 hardware.*

```
dm v1.16
console v1.0
        Copyright Alliant
Modules In System:
  mm0a:32M mm0b:32M
     |
    cm0:.5M
      |
dsw0 dsw1 dsw2 dsw3 dsw4 dsw5 dsw6 dsw7 dsw8 dsw9 dswa dswb dswc dswd dswe dswf
 asw0 asw1 asw2 asw3 asw4 asw5 asw6 asw7 asw8 asw9  msw  clk
             | |                          | |
            iom1                         pm4
            sce4                        sce10
                                        sce11
                                        sce12
                                        sce13
Turning cm0 A B and C Status Latches OFF!
Total run time (in hours): 0.000 Total Errors: 0
```

*The FX/2800 booter loads Concentrix-2800 under debugger control.*

```
Alliant FX/2800 Booter
/dev/rxdv00r:/ddt -> 0x1b000+0xc000+0x1e418, syms = 0x42c0, strings = 0x2f84
/dev/rxdv00r:/unix -> 0xf2000+0x20000+0x3f668, syms = 0x10580, strings = 0xc555
IOM Started.
```

*The "tip" program connects the console to the Concentrix-2800 system.*

```
tip escape character is ^^ (control-shift-6)
connected
ddt initializing; state FFFFEDF0 from FFFFEF90
Alliant ddt 6.0 built Sun May 20 20:48:27 EDT 1990.
Built by root (/src1/Sunrise/sr1src/usr/src/sys/srddt)
```

*The system stops in Alliant ddt 2800 mode if the d flag is in effect; the debugger does not prompt. Type* ESCAPE *p or* ESCAPE *g to proceed from ddt.*

ESCAPE *g*

```
ip00>
*-------------------------------------------------------------------*
|   Alliant Concentrix 2800 2.1.00 [built] Mon Jan 21 12:56:08 EST 1991 |
|              root (/rel1/sr2.1src/usr/src/sys/CONCENTRIX)          |
| Copyright (c) 1985-1990 Alliant Computer Systems Corp. (Unpublished) |
|                        All Rights Reserved                        |
*-------------------------------------------------------------------*
ip00>Real Memory: 16384 pages, 64.0 Mb
ip00>Cache Memory: 128 pages, 512 Kb
ip00>Palloc Memory: 714 pages, 2.8 Mb
ip00>124 base system locks
ip00>1358 lock copies (1482 total system locks)
ip00>0 statistics avail., 0 trace buckets avail.
ip00>Available Memory: 9611 pages, 37.5 Mb
ip00>Using 1638 buffers containing: 1638 pages, 6.4 Mb
ip00>Configured for 64 users
ip00>Configuring IP0 (0x0)
ip00>console stream[0]: configured
ip00>console stream[1]: configured
ip00>*** Attaching VME-AC0 devices:
ip00>   installing root: ramdisk00r
ip00>ramdisk: "ram4mb" <823.,4.,512.> spares=0
ip00>xdv0**: 754, prom=2282, firmware=2.4.0
ip00>   xdv00*: "fj1000s" <745,27,80+2> [bsize=512]
ip00>   installing swap: xdv00s
ip00>   xdv01*: "fj1000s" <745,27,80+2> [bsize=512]
ip00>vme_scsi@0xfd00c000: code(077,0) firmware(XAL,03/22/90)
ip00>   exabyte[0,2]: EXABYTE EXB-8200        425A sequential (removable)
ip00>ex[0]: online, versions: nx 5.5, ex 2.0
ip00>mub_init: IP0. max iomaps=1936. 7740Kb transfers possible.
ip00>1024 pages available for user ISRs at address 0xfc800000
ip00>Configuring 1024 refbuf, 1 avgbuf, 2 maxbuf
ip00>Configuring 1024 refino, 1 avgino, 1596 totino
ip00>Configuring 64 refsno, 2 avgsno, 128 totsno
ip00>Configuring 512 refrno, 2 avgrno, 128 inirno, 1024 maxrno
ce00>Configuring CE0 (0x8)
ce01>Configuring CE1 (0xa)
ce02>Configuring CE2 (0x9)
ce03>Configuring CE3 (0xb)
ce04>Configuring CE4 (0xc)
ce05>Configuring CE5 (0xe)
ce06>Configuring CE6 (0xd)
ce07>Configuring CE7 (0xf)
ip00>DETACHED: CE0 CE1 CE2 CE3 CE4 CE5 CE6 CE7
                IP0 [9 CE/IP's]
ip00>Paging File space: 42322 pages, 165.3 Mb
ip00>Swapping File space: 34627 pages, 135.3 Mb
ip00>Available user memory: 9419 pages, 36.8 Mb
```

*The system stops in single-user mode if the s flag is in effect. Type* CTRL d *to continue.*

```
Concentrix-sh# CTRL d
```

*Processors can be attached/detached by: specifying the clcr and clrm commands; adjusting /etc/console.config; or (if /etc/console.config is not used) shutting down to Diagnostix-2800 (shutdown -h), adjusting /console/console.config, and rebooting. The console.config files can also be adjusted to kill (not configure) hardware elements.*

*The standard start-up procedure /etc/rc executes. Its first operation is usually to check the consistency of the file system using fsck. Fsck may produce a long series of messages as it repairs inconsistencies. This example shows a fast boot – the file system is not checked.*

```
Concentrix-sh# Fast boot ... skipping disk checks
ip00> ex[0]: [8:0:14:21:28:65]
```

*/etc/rc runs the site-specific script /etc/rc.local at this point. /etc/rc.local starts the software required on this system and can be tailored by the system administrator.*

```
Starting rpc services: portmap.
starting system logger

starting local daemons: batch.

Starting rtnamed
 (logical name table "tmp" created).
preserving editor files
clearing /tmp
standard daemons: update cron accounting.
starting network daemons: rwhod inetd printer.
Thu Jul 12 19:02:29 EDT 1990
```

*/etc/rc.local finishes and returns control to /etc/rc. /etc/rc executes the last few operations required to start multi-user Concentrix-2800.*

```
ALLIANT CONCENTRIX (duane-sr)

login:
```

*At this point, you can log into Concentrix-2800 on the console.*

## 1.4.5 System Initialization

The system start-up procedure creates and passes control to a process named init (/etc/init). Init is the parent of all other processes and has a process identification (pid) of 1. See init in Appendix A for a description.

Init executes a start-up script depending on the system being started. If Diagnostix-2800 is being started, the script is /etc/rc.diag. If Concentrix-2800 is being started, the script is /etc/rc, which in turn executes /etc/rc.local.

### 1.4.5.1 Script rc.diag

The script /etc/rc.diag as distributed is reproduced below. You must be connected to Diagnostix-2800 to access the script; the script should need no modifications.

```
if [ -r /fastboot ]
then
        rm -f /fastboot
        echo fastboot... skipping disk checks
elif [ $1x = autobootx ]
then
        date
        echo checking root partition...
#
#       To check the root device, delete the '#' preceding the
appropriate
#       commmand line.
#
        /etc/fsck -p /dev/xdv00p
        case $? in
        0)
                date
                ;;
        4)
                /etc/reboot -n -`/console/bootflags`
                ;;
        8)
                echo "automatic reboot failed... help!"
                exit 1
                ;;
        12)
                echo "reboot interrupted"
                exit 1
                ;;
        *)
                echo "unknown error in reboot"
                exit 1
                ;;
        esac
else
        date
fi

/etc/mount -f /dev/xdv00p /
echo "FX/2800 Startup"

:
: Start Concentrix
:

/console/console_loop $1

: return to single user mode
kill -TERM 1
```

### 1.4.5.2 Script rc

The script /etc/rc as distributed is reproduced below. If your system uses file system quotas, uncomment the four lines for checking and turning quotas on. If your system is *not* on a network, comment out the three lines near the end of the file that start rwhod.

```
bootflags=$2

HOME=/; export HOME
PATH=/bin:/usr/bin

/etc/mkstripe -a>/dev/console
if [ -r /fastboot ]
then
        rm -f /fastboot
        echo Fast boot ... skipping disk checks >/dev/console
elif [ $1x = autobootx ]
then
        echo Automatic reboot in progress... >/dev/console
        date >/dev/console
        /etc/fsck -p >/dev/console
        case $? in
        0)
                date >/dev/console
                ;;
        2)
                exit 1
                ;;
        4)
                /etc/reboot -n $bootflags
                ;;
        8)
                echo "Automatic reboot failed... help!" >/dev/console
                exit 1
                ;;
        12)
                echo "Reboot interrupted" >/dev/console
                exit 1
                ;;
        *)
                echo "Unknown error in reboot" > /dev/console
                exit 1
                ;;
        esac
else
        date >/dev/console
fi

rm -f /dev/swap*
/bin/mkswap
```

```
# attempt to rationally recover the passwd file if needed
if [ -s /etc/ptmp ]
then
        if [ -s /etc/passwd ]
        then
                ls -l /etc/passwd /etc/ptmp >/dev/console
                rm -f /etc/ptmp         # should really remove the shorter
        else
                echo 'passwd file recovered from ptmp' >/dev/console
                mv /etc/ptmp /etc/passwd
        fi
elif [ -r /etc/ptmp ]
then
        echo 'removing passwd lock file' >/dev/console
        rm -f /etc/ptmp
fi

/etc/umount -a
cp /dev/null /etc/mtab
#/etc/swapon -a                                          >/dev/console 2>&1
/etc/mount -f /                                         >/dev/console 2>&1
/etc/mount -at 4.2                                      >/dev/console 2>&1
/etc/mount -at tfs                                      >/dev/console 2>&1

# Handle quotas (Remove leading #'s from each of the four lines
#                following the next blank line when activating
#                quotas for your system.)

#                               echo -n 'checking quotas: '     >/dev/console
#/etc/quotacheck -a -p                                          >/dev/console
2>&1
#                               echo  'done.'                   >/dev/console
#/etc/quotaon -a                                                >/dev/console
2>&1

rm -f /etc/nologin
rm -f /usr/spool/uucp/LCK.*
chmod 666 /dev/tty[pqrs]*

sh /etc/rc.local
                              echo preserving editor files      >/dev/console
(cd /tmp; /usr/lib/ex3.7preserve -a)

if [ -d /tmp ]
then
                              echo clearing /tmp                >/dev/console
(cd /tmp; find . ! -name . ! -name lost+found ! -name quotas -exec rm -r {} \; )
else
     echo /tmp directory does not exist, cannot clear /tmp      >/dev/console
fi


                              echo -n standard daemons:         >/dev/console
/etc/update;                  echo -n ' update'                 >/dev/console
/etc/cron;                    echo -n ' cron'                   >/dev/console
/etc/accton /usr/adm/acct;    echo -n ' accounting'             >/dev/console
                              echo '.'                  >/dev/console
```

```
echo -n starting network daemons:                                    >/dev/console
if [ -f /etc/rwhod ]; then
        /etc/rwhod;             echo -n ' rwhod'        >/dev/console
fi
if [ -f /etc/inetd ]; then
        /etc/inetd;             echo -n ' inetd'        >/dev/console
fi
if [ -f /usr/lib/lpd ]; then
        rm -f /dev/printer
        /usr/lib/lpd;           echo -n ' printer'              >/dev/console
fi
                                echo '.'                >/dev/console
                                date                            >/dev/console
#
# start up graphics daemons (X11, PHIGS...)
#
if [ -f /etc/rc.graphics ]; then
        sh /etc/rc.graphics; echo -n 'start up graphics daemons'  >/dev/console
fi
echo '.'                                                >/dev/console
exit 0
```

### 1.4.5.3 Script rc.local

The script /etc/rc.local as distributed is reproduced below. Typically, you edit this script during system installation. Comments describe the edits that should be made depending on your site requirements.

```
PATH=/etc:/usr/etc:/bin:/usr/ucb:/usr/bin:/usr/alliant

# myname is my symbolic name
# myhostid is my internet address in #.#.#.# format
# my-netmask is specified in /etc/networks
#
```

*Replace* myname *with the name of your system.*

```
hostname myname
```

*If you are part of a network information service (NIS) domain, uncomment the next line and replace* mydomain *with the name of the domain.*

```
#domainname mydomain
```

*(Use the commented out version of* ifconfig *if you have an old Excelan controller. Uncomment the* ex0 *version and comment out the* en0 *version.) Modify the first ifconfig line to suit your network configuration. If you are not running subnets, you can delete* netmask my-netmask *and* broacast my-broadcast. *At any rate, ensure that* my-broadcast *identifies your network, not your host (the last decimal position must be 0). Delete this line if you have no Ethernet controllers. If you have more than one Ethernet controller, specify additional ifconfig lines. See ifconfig in Appendix A. Do not modify the second ifconfig line (for localhost).*

```
#ifconfig ex0 -trailers netmask my-netmask broadcast my-broadcast `hostname`
ifconfig en0 -trailers netmask my-netmask broadcast my-broadcast `hostname`
ifconfig lo0 localhost
sleep 3
route add `hostname` localhost 0
hostid `hostname`
```

```
echo -n 'Starting rpc services:'                        >/dev/console
if [ -f /etc/portmap ]; then
        portmap; echo -n ' portmap'                     >/dev/console
fi
```

*Uncomment the next 3 lines if you are an NIS server.*

```
#if [ -f /etc/ypserv -a -d /etc/yp/`domainname` ]; then
#       ypserv; echo -n ' ypserv'                       >/dev/console
#fi
```

*Uncomment the next 3 lines if you are an NIS slave.*

```
#if  [ -f /etc/ypbind ]; then
#       ypbind; echo -n ' ypbind'                       >/dev/console
#fi
```

*Uncomment the next 3 lines for the network routing daemon.*

```
#if [ -f /etc/routed ]; then
#       /etc/routed;            echo -n ' routed'       >/dev/console
#fi

echo '.'                                                >/dev/console
mount -vat nfs                                          >/dev/console
adb /vmunix - << EOF |  sed -e "s/._versionname://" >/tmp/t1
versionname?s
EOF
tail +2 /etc/motd | grep -v "Concentrix 2800" >>/tmp/t1
mv /tmp/t1 /etc/motd
chmod 666 /etc/motd

# syslogd doesn't belong here, but needs to be started before the others.
if [ -f /etc/syslogd ]; then
                        echo 'starting system logger' >/dev/console
        rm -f /dev/log
        /etc/syslogd
fi

#               echo -n 'checking for core dump... '    >/dev/console
#/usr/etc/savecore /usr/crash                           >/dev/console 2>&1
                echo ''                                 >/dev/console
```

*Uncomment the next 3 lines for the Internet domain name server.*

```
                        echo -n starting local daemons:         >/dev/console
#if [ -f /etc/named ]; then
#       /etc/named /etc/named.boot; echo -n ' named' >/dev/console
#fi
```

*Uncomment the next 3 lines for the time server daemon.*

```
#if [ -f /etc/timed ]; then
#       /etc/timed -M;          echo -n ' timed'        >/dev/console
#fi
if [ -f /usr/lib/sendmail -a -f /usr/lib/sendmail.cf ]; then
        (cd /usr/spool/mqueue; rm -f lf*)
        /usr/lib/sendmail -bd -q30m; echo -n ' sendmail'>/dev/console
fi
if [ -f /usr/lib/batch/batchd ]; then
        rm -f /usr/spool/batch/qtmp/*
        /usr/lib/batch/batchd;
        (cd /usr/lib/batch; ./netdae -n 2 /usr/lib/batch/netrecv) &
        echo -n ' batch'                                >/dev/console
fi
```

*Uncomment the next 8 lines for NFS. Specify "/etc/nfsd 8" instead of "/etc/nfsd4" if you are running diskless/NFS.*

```
#if [ -f /etc/nfsd -a -f /etc/exports ]; then
#       > /etc/xtab
#       /usr/etc/exportfs -a
#       /etc/nfsd 4 & echo -n ' nfsd'                    >/dev/console
#fi
#if [ -f /etc/biod ]; then
#       /etc/biod 4; echo -n ' biod'                     >/dev/console
#fi
```

*Uncomment the next 5 lines for diskless/NFS.*

```
#if [ -d /tftpboot ]; then
#       (echo -n ' bootparamd'; \
#       /usr/etc/rpc.bootparamd)                         >/dev/console
#       arp -e /etc/ethers                               >/dev/console
#fi

#
#
#
# start up status monitor and locking daemon if present
#
```

*Uncomment the next 3 lines for the network status monitor.*

```
#if [ -f /etc/rpc.statd ]; then
#       /etc/rpc.statd & echo -n ' statd'            >/dev/console
#fi
```

*Uncomment the next 3 lines for the network lock daemon.*

```
#if [ -f /etc/rpc.lockd ]; then
#       /etc/rpc.lockd & echo -n ' lockd'            >/dev/console
#fi
#
#start up the rtname daemon
#
```

*Uncomment the next 10 lines for logical name tables.*

```
#if [ -f /usr/alliant/rtnamed ]; then
#       echo ''                                      >/dev/console
#       /usr/alliant/rtnamed & echo -n ' rtnamed'    >/dev/console
#
#   if [ -f /usr/alliant/rtncrtbl ]; then
#         sleep 3
#       /usr/alliant/rtncrtbl tmp 128 0 -w:elr -g:elr
#       echo -n ' (logical name table "tmp" created)'   >/dev/console
#   fi
#fi
# initialize gosip
if [ -f /etc/rc.gosip ]; then
        sh /etc/rc.gosip ;   /bin/echo 'initializing gosip' >/dev/console
fi
#
#
# Set up complex
#
if [ -f /etc/rc.cluster ]; then
    .   echo ' configure cluster'                    /dev/console
       csh /etc//rc.cluster
fi
echo '.'                                             >/dev/console
```

The following example demonstrates the first 13 lines of /etc/rc.local for the case where all the systems in the network are running the same level of network code (that is, all systems are running a 4.3BSD level of network code).

```
PATH=/etc:/usr/etc:/bin:/usr/ucb:/usr/bin:/usr/alliant

# myname is my symbolic name
# myhostid is my internet address in #.#.#.# format
# my-netmask is specified in /etc/networks
#
hostname odie
#domainname mydomain
#ifconfig ex0 -trailers netmask my-netmask broadcast my-broadcast `hostname`
ifconfig en0 -trailers `hostname`
ifconfig lo0 localhost
sleep 3
route add `hostname` localhost 0
hostid `hostname`
```

In the second example, the system is in a mixed 4.2BSD/4.3BSD environment, with NFS and NIS running.

```
PATH=/etc:/usr/etc:/bin:/usr/ucb:/usr/bin:/usr/alliant

# myname is my symbolic name
# myhostid is my internet address in #.#.#.# format
# my-netmask is specified in /etc/networks
#
hostname odie
domainname onet
#ifconfig ex0 -trailers netmask my-netmask broadcast my-broadcast `hostname`
ifconfig en0 -trailers broadcast 125.0.0.0 `hostname`
ifconfig lo0 localhost
sleep 3
route add `hostname` localhost 0
hostid `hostname`
```

## 1.4.6 Boot from RAM Disk

A RAM disk simulates a disk device with physical memory. You can start Concentrix-2800 in single-user mode using a RAM disk for the root file system, for example, if your root disk fails. The supplied RAM disk root file system is limited to disk maintenance and install utilities; see Appendix D for more information.

To boot from a RAM disk, choose option R at the console boot loop prompt (FX/2800 startup mode). Booting from a RAM disk is disabled by default; enter e to enable, then reboot by responding y:

```
Reboot with flags set to ""? (y for yes, ? for help):R

Ramdisk booting is currently disabled.
Do you want to enable (e) or disable (d) ?e

Reboot with flags set to ""? (y for yes, ? for help): y
```

The system is now set to boot from /console/ramunix and /console/ramddt, and to start srboot with the switch -r ramdisk.image, which loads /console/ramdisk.image into physical memory, and arranges to use this image as the initial root file system.

During loading the following messages may appear:

```
Turning cm4 A B and C Status Latches OFF!
Total run time (in hours): 0.000 Total Errors: 0
Booting from ramdisk ***
Alliant FX/2800 Booter
Allocating 3072 pages for ramdisk image of 1024 pages. ***
Ramdisk of 3072 pages loaded with 4194304 bytes. ***
ramddt -> 0x1f000+0x10000+0x2b01c, syms = 0x44e0, strings = 0x3137
ramunix -> 0x107000+0x21000+0x44b74, syms = 0x10ea0, strings = 0xcdac
IOM Started.
tip escape character is ^^ (control-shift-6)
connected
```

Then the startup banner notifies you that ramdisk is being used as the root (lines marked with stars highlight these lines):

```
ip00> Using 819 buffers containing: 819 pages, 3.2 Mb
ip00> Configured for 64 users
ip00> Configuring IP0 (0x4)
ip00> console stream[0]: configured
ip00> console stream[1]: configured
ip00> bus 0, intr @ level 3, vec 0x43 acked
ip00> *** Attaching VME-AC0 devices:
ip00>     installing root: ramdisk00r                ******
ip00> ramdisk: "ram4mb" <823.,4.,512.> spares=0      ******
ip00> xdv0**: 754, prom=2282, firmware=2.4.0
ip00>     xdv00*: "fj690s" <624,27,64+4> [bsize=512]
ip00>     installing swap: xdv00s
ip00> vme_scsi@0xfd00c000: code(091,0) firmware(XAN,06/26/90)
```

Then the Concentrix Utilities Menu appears; see Appendix D.

The flag to boot from the RAM disk only works once. It must be reenabled if you want to use it again.

See Appendix A, makeramdisk for a description of how to create a non-root file system on a RAM disk for use in multi-user Concentrix-2800.

## 1.5 DATE AND TIME

The date command typed without an argument displays the current day, date, time, and time zone. The following example shows the current day as Wednesday, the current date as March 22, 1990, the current time as 6:39 P.M. and 24 seconds, and the time zone as Eastern Standard Time:

```
# date
Thurs Mar 22 18:39:24 EST 1990
```

If the time zone on your machine is not correct, shut down to single-user mode and enter the correct time zone with the command /usr/alliant/timezone. The following example resets the time zone to Pacific U.S. time:

```
# /usr/alliant/timezone pt us
```

Then change directory to /etc/zoneinfo, remove the file localtime (which is delivered as a soft link to /etc/zoneinfo/EST5EDT) and re-create it as a soft link to the all-uppercase file in /etc/zoneinfo that describes your time zone. The example resets localtime for Pacific time:

```
# cd /etc/timezone
# rm localtime
# ls -s PST8PDT localtime
```

Then type CTRL D to return to multi-user Concentrix-2800.

If the date or time is incorrect (correct the time zone first), type the date command with an argument that specifies the correct date and time. The following example sets the date to 1990, March (month 03) 22, 6:39 P.M. (1839):

```
# date 9003221839
```

## 1.6 SYSTEM FAILURES

You should notify Alliant Customer Service of a system failure. Use one of the following telephone numbers:

- **1-800-453-5305** – Inside Massachusetts, toll-free.
- **1-800-225-4268** – Outside Massachusetts, inside United States, toll-free.
- **508-486-4950** – Outside United States, not toll-free.

Call_cs in Appendix B lists the telephone numbers and additional information that you should have on hand when you call.

During normal working hours, you should await instructions from Alliant Customer Service before attempting to restart the system. Alliant Customer Service may want to dial into the system and diagnose the problem.

During off hours, you may want to try to restart the system yourself. In either case, log the problem and any actions that are taken.

### 1.6.1 Terminal Hang

If only one terminal does not respond, the problem is probably the terminal or the applications program running on the terminal. (A program hang may also cause a continuous stream of typing to the terminal, for example, if the program hangs in a loop that types to standard output.) Try the following remedies:

- **Terminal** – Press CTRL Q to make sure the terminal is not simply in stop mode. Power the terminal off and on.
- **Program** – Press CTRL C to kill the program. If CTRL C does not kill the program, press CTRL Z to stop the job under which the program is running; type jobs RETURN to get the number of the job; type kill –9 *%number* RETURN.

If none of the above remedies work, log in on another terminal as root and identify the processes of the user of the hung terminal, usually by piping the output of ps aux through grep. Then kill the offending process; the second column of the ps aux output contains the process identification (pid). Note the following example:

```
% ps aux | grep jones
jones     1122   284 L   168 3a Slp     0:05  -csh  (csh)
root      3394    56 L    64 14 Slp     0:00  grep jones
jones     3397    80 L    56 3a Run     0:02  test221
% kill -9 3397
```

If the terminal remains hung, kill the login shell (job 1122 in the example). If the terminal still remains hung, call Alliant Customer Service. You may have to restart the system (fastboot) to recover the terminal.

## 1.6.2 System Hang

A system is hung when all the terminals stop responding. The cause of a hang could be hardware or software. Use one of the following procedures to restart Concentrix-2800:

- Connect to FX/2800 startup mode by typing CTRL^ (hold down the control and shift keys, then press the 6 key) followed by a period. From this mode, you can type y or other options to restart Concentrix-2800.
- Connect to Diagnostix-2800 mode by typing CTRL^ (hold down the control and shift keys, then press the 6 key) followed by an exclamation mark. From this mode, you can type CTRLd to enter FX/2800 startup mode, or type a reboot or halt command.
- Enter ddt, the kernel debugger, by typing CTRL^ (hold down the control and shift keys, then press the 6 key) followed by a pound sign. Perform one of the preceeding procedures when you are finished in the debugger.

See the *Concentrix-2800 Systems Programming Guide* for a description of the kernel debugger. Also, check the procedure Recover_system in Appendix B.

## 1.6.3 System Panic

A system panic is obvious because the system types information on the console and invokes ddt, the kernel debugger. Barring other instructions from Alliant, you should first dump memory as described under Recover_system (Appendix B), then connect to FX/2800 startup mode or Diagnostix-2800 mode to reboot the system.

A panic is the result of the failure of an internal consistency test. The message typed to the console has the following form. The panic string is typically the name of the routine that detected the error, or a short description of the inconsistency.

```
panic: panic string
```

The most common panic strings are as follows:

```
realloccg: bad optim
ialloc: dup alloc
alloccgblk: cyl groups corrupted
ialloccg: map corrupted
free: freeing free block
free: freeing free frag
ifree: freeing free inode
```

These panic messages occur when file system inconsistencies are detected. The problem generally results from a failure to repair damaged file systems after a crash, hardware failures, or other conditions that should not normally occur. A file system check normally corrects the problem.

`rootmount cannot mount root`

This message results from a failure to mount the root file system during the boot process. Either the root file system is corrupted, or the system is attempting to use the wrong device as the root file system.

`kmem_alloc`

This error occurs if the system cannot handle an internal memory request. Increasing the kernel variable pallocsize or vmepallocsize should fix the problem.

`LCK_exclusive: hierarchy violation`

The symmetric multi-processor design of Concentrix-2800 requires multiple locks to serialize access to data structures, and a hierarchy prevents deadlocks among processors holding multiple locks. This panic means that the hierarchy protocol was violated.

`Can't exec /etc/init`

Late in the boot procedure, the system is unable to locate and execute the initialization process init, the root filesystem is incorrect or corrupted, or the mode or type of /etc/init forbids execution.

## 1.6.4 Peripheral Errors

When errors occur on peripheral devices, the system prints a warning diagnostic on the console. Device error messages are documented in the *Concentrix System Reference: Miscellaneous, Special Files, and File Formats*. All system messages are collected regularly and written into the system error log file /usr/adm/messages. You should examine the error log file regularly (tail –r /usr/adm/messages).

# CHAPTER 2
# USER ACCOUNTS

The following tables outline the commands, procedures, and files pertaining to user accounts.

**Administration Files**

| name | description |
|---|---|
| ~ | user login directory |
| /quotas | quota file at file system root (edit) |
| ~/.hushlogin | file that inhibits message of the day |
| ~/.forward | mail forwarding (aliasing) file |
| ~/.mailrc | mail start-up file |
| ~/.msgsrc | next message to be processed by msgs command |
| ~/.plan | user-supplied text displayed by the finger command |
| ~/.project | user-supplied text displayed by the finger command |
| /etc/gettytab | terminal configuration data base ** (edit) |
| /etc/group | group file * (edit) |
| /etc/motd | message of the day (edit) |
| /etc/nologin | file whose existence prevents logins |
| /etc/passwd | password file * (edit) |
| /etc/passwd.dir | password file - data base format |
| /etc/passwd.pag | password file - data base format |
| /etc/ptmp | temporary file |
| /etc/syslog.conf | syslog configuration file |
| /etc/syslog.pid | syslog process identification |
| /etc/ttys | terminal initialization data * (edit) |
| /etc/utmp | login record of current users** |
| /usr/adm/acct | execution accounting file * |
| /usr/adm/savacct | summary of execution accounting file |
| /usr/adm/usracct | per-user summary of execution accounting file |
| /usr/adm/wtmp | login record of all logins and logouts * |
| /usr/lib/aliases | aliases file for sendmail ** (edit) |
| /usr/lib/aliases.dir | binary representation of aliases file |
| /usr/lib/aliases.pag | binary representation of aliases file |
| /usr/lib/Mail.rc | system-wide mail start-up file (edit) |
| /usr/lib/sendmail.cf | configuration file for sendmail |

**continued on next page**

**Administration Files (continued)**

| name | description |
|---|---|
| /usr/lib/sendmail.hf | help file for sendmail |
| /usr/msgs/* | messages data base |
| /usr/skel/.??* | skeleton files |
| /usr/spool/mqueue/* | sendmail queue and system log |
| /usr/spool/mail/* | mail queue |
| /usr/spool/secretmail/* | encrypted mail queue |

* See *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

**Administration Commands (Appendix A)**

| name | description |
|---|---|
| ac | login accounting |
| accton | enable/disable system accounting |
| chown | change owner |
| comsat | biff server |
| edquota | edit user quotas |
| getty | set terminal mode |
| mailq | print contents of queue (sendmail) |
| mkpasswd | generate hashed password table |
| newaliases | rebuild data base for mail aliases file |
| quotacheck | file system quota consistency checker |
| quotaon/quotaoff | turn file system quotas on and off |
| repquota | summarize quotas for a file system |
| sa | maintain system accounting files |
| sendmail | send mail over the internet |
| syslogd | log system messages |
| vipw | edit the password file |

**Administration Procedures (Appendix B)**

| name | description |
|---|---|
| Add_user | add a new user to the system |
| Add_quotas | add file system quotas |
| Broadcast | broadcast messages to users |
| Install_mail | install mail on a new system |
| Remove_user | remove a user from the system |

**General Commands (Concentrix-2800 Commands and Applications Manual)**

| name | description |
|---|---|
| biff | be notified if mail arrives |
| chgrp | change file group |
| chkey | change your encryption key |
| chmod | change file access mode |
| keylogin | decrypt and store secret key |
| login | create new user process |
| logout | terminate login shell |
| /usr/ucb/mail | send and receive mail |
| /bin/mail | send and receive mail (old) |
| mkdir | create directory |
| msgs | read system messages |
| passwd | change login passwd |
| passwd -f | create/change finger entry |
| quota | display disk usage and limits |
| rlogin | remote login |
| su | substitute userid (superuser) |
| telnet | user interface to the TELNET protocol |
| wall | write to all users |
| xget | read encrypted mail |
| xsend | send encrypted mail |

## 2.1 LOGIN PROCEDURE

A user gains interactive access to Concentrix-2800 by pressing RETURN on a terminal that is connected to the system. (The user may have to dial in through a modem or observe some other protocol to connect the terminal to the system.) A getty process then prompts for and collects the user's name, and invokes a login process. The login process verifies the user password, initializes the user's environment, performs other activities, and leaves the user in a shell (as specified in the /etc/passwd file). See getty and login in Appendix A for details.

A user can terminate interactive access to Concentrix-2800 with the logout command. A user can log out and log in again with the login command. The rlogin and telnet commands permit interactive access to other systems (or the same system) while retaining access to the local system.

Use the su command to temporarily change your user identification without logging out. Normally, you must know the password of the new user identification. If you are in the wheel group, typing su without an argument changes your user identification to root. If you are superuser, typing su *userid* changes your user id to *userid*.

After attaining the new user identification, you must type cd (places you in the new user's login directory) and source .login (executes the new user's login file) to simulate an actual login. Type exit to terminate the substitute process and return to your original process.

## 2.2 PASSWD FILE

The password files /etc/passwd (source), /etc/passwd.dir (data base), and /etc/passwd.pag (data base) identify the users that can log into the system. All three files are needed for correct operation of the system. Absence of a data base file typically results in a dbm_open failed message.

The file /etc/passwd identifies each user that can log into the system. The file is a plain ASCII file; each line represents one user entry (that is, user entries are separated by newline characters). Each user entry contains fields, which must be separated by colons; two consecutive colons indicate that the contents of a field are not specified. The specification of a user entry is as follows:

*name:password:user-id:group-id:gcos-info:home:shell*

The Add_user procedure in Appendix B provides more detail on the fields as does the passwd entry in the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats.*

The passwd file should be edited with the /etc/vipw command (see Appendix A), which performs the following services:

- Locks /etc/passwd while the edit is in progress. In addition, the lock file /etc/ptmp serves as a backup in case the system crashes while the edit is in progress; recovery is part of the start-up script rc.
- Permits you to edit the file with the editor defined by the environment variable EDITOR (defaults to vi).
- Executes /etc/mkpasswd to create the data base password files /etc/passwd.dir and /etc/passwd.pag. The data base files are necessary for some operations and must be made after an update to /etc/passwd.

Typically, you provide a name, user identification, group identification, the GCOS information, a home directory, and a shell (defaults to the Bourne shell). Use the passwd command to fill in the *password* field. Use the chfn command to replace the contents of the *gcos* field. The following example demonstrates an /etc/passwd entry for a user named hacker, whose full name is Polly Hacker as indicated in the *gcos* field:

```
hacker::301:300:Polly Hacker:/staff/usr/hacker:/bin/csh
```

The user name (for example, hacker) and user identification (for example, 301) must be unique.

The /etc/passwd file that you receive with a new system has the following contents:

```
root::0:10:Charlie &:/:/bin/csh
daemon:*:1:31:The devil himself:/:
batch:*:2:1:batch:/:
sys:*:2:2:/:/bin/csh:
bin:*:3:3:/bin::
UUCP::66:10:UNIX-to-UNIX Copy:/usr/spool/uucppublic:/usr/lib/uucp/uucico
uucp::66:10:UNIX-to-UNIX Copy:/usr/spool/uucppublic:/usr/lib/uucp/uucico
duu::11:11:dialup-user:/:/bin/true
cs::900:10:Alliant Customer Service:/usr/cs:/bin/csh
nobody:*:65534:65534::/:
```

You should log into (or substitute user to) each account in the file and provide passwords.

If you are using the Network Information Service, see Chapter 7.

## 2.3 GROUP FILE

The file /etc/group associates names with group identifications and also identifies the users who belong to the groups. Membership in a group provides group access to files owned by members of the group. A user can belong to a maximum of 16 groups. (Note, however, that Sun workstations only support 8 groups per user.)

A user belongs to one primary group, which is specified by identification number (not name) in field 4 of the user's record in /etc/passwd.

See the commands chmod (*Concentrix-2800 Commands and Applications Manual*), chown (Appendix A), and chgrp (*Concentrix-2800 Commands and Applications Manual*) for further information on file ownership and access restrictions.

Group ownership for new files is determined by the group ownership of the parent directory (not the primary group of the creator of the file). The system administrator must be sure to set the correct group ownership when creating user directories.

The group file is a plain ASCII file; each line represents one user entry (that is, user entries are separated by newline characters). Each user entry contains fields, which must be separated by colons; two consecutive fields indicate that the contents of a field are not specified.

*group-name*:*password*:*group-id*:*user-id*[,*user-id*,...]

The Add_user procedure in Appendix B provides more detail on the fields as does the group entry in the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*.

The group file is edited with any editor. Typically, you provide a name for the group, the group identification number, and a comma-separated list of user names. The *passwd* field is typically not specified. The following example demonstrates an /etc/group entry for group 300:

```
algorithms::300:able,baker,charlie,dog,echo,fox,gorn,hacker
```

The /etc/group file that you receive with a new system has the following contents:

```
wheel:*:0:root
daemon:*:1:daemon
kmem:*:2:root
sys:*:3:root
tty:*:4:root
staff:*:10:root
operator:*:28:root
```

```
guest:*:31:root
ingres:*:74:ingres
realtime:*:99:root
```

If you are using the Network Information Service, see Chapter 7.

## 2.4 SKELETON FILES

In adding a new user to the system, you typically create a home directory for the user and copy to that directory the files contained in the directory /usr/skel. The procedure Add_user in Appendix B provides the detailed steps. The skeleton files are as follows:

- .cshrc – A script that is invoked at the start of a new C shell, for example, when a user logs in and the user's login shell as specified in /etc/passwd is /bin/csh.
- .login – A script that is invoked when a user logs in (after .cshrc).
- .logout – A script that is invoked when a user logs out.
- .mailrc – A script that is invoked as part of the mail start-up procedure.

The user may want to modify these files as experience on the system advances. Similarly, the system administrator may want to modify the skeleton files themselves. For example, the umask command (set file creation mask) in .login is 002, which provides default file permissions of 775 (read and write access for owner and group, read access for all); you may want to change the mask to 022, which provides permissions of 755 (group cannot write), or 077, which provides permissions of 700 (no access except for owner).

You may also want to provide new users with information files used by the finger command (see the *Concentrix-2800 Commands and Applications Manual*). These files are also placed in the user's home directory:

- .plan – Typically contains information about what the user is working on, but can contain anything.
- .project – Typically contains information about what the user is working on, but can contain anything. Restricted to 70 characters.

## 2.5 USER COMMUNICATIONS

The following facilities enable the system administrator to broadcast information to all users. See Broadcast in Appendix B for details.

- **Message of the day** – Typically, the system administrator keeps a short list of critical messages in a plain text file named /etc/motd. The file is displayed on a user's terminal at login (unless the user's login directory contains a file named .hushlogin). You should delete a message from the file as soon as it is no longer needed.

- **System messages** – The system administrator can collect messages in the directory /usr/msgs through the msgs –s command or by mailing messages to user msgs. The messages can be accessed by users through the msgs command. Typically, each user's .login file contains the msgs command.

- **Write to all users** – The system administrator can immediately broadcast a message to every user terminal with the wall command.

## 2.6 MAIL

The mail facility collects messages from users and transfers them to the sendmail facility, and provides users access to messages in the local mail queue. The sendmail facility routes internetwork messages to sendmail facilities on other systems and routes local messages to the local mail facility, which places them in the local mail queue.

The mail facility must be set up (one time) on a new system. See Install_mail in Appendix B for details.

### 2.6.1 Mail Command

The system has two mail commands: the newer, preferred version has the pathname /usr/ucb/mail; the older version has the pathname /bin/mail. Users should place the appropriate directory (/usr/ucb or /bin) first in their command paths, depending on the desired version. The skeleton file /usr/skel/.login places /usr/ucb first:

```
set path = ( . $path[1] ~/bin /usr/local $path[2-] /usr/alliant /usr/sys5 )
```

You can place /bin first by changing $path[2-] to /bin /usr/ucb /usr/bin.

Alternatively, you can place an alias in the skeleton file /usr/skel/.cshrc or users can place an alias in their individual .cshrc files (applies to C shell users):

```
alias mail /bin/mail
```

### 2.6.2 Mail Queue (User Mailboxes)

The mail queues are files in the directory /usr/spool/mail. Each file is named after the user whose mail is there; the file is owned by the user with read and write permissions for the owner only.

### 2.6.3 Start-up Files

The file /usr/lib/Mail.rc, if it exists, sets mail options at the start of each mail session by any user on the system. See the *Concentrix System User's Guide*, "Mail Reference," for a list of options. Note the following example:

```
% cat /usr/lib/Mail.rc
set append ask askcc
set crt=24 hold
set EDITOR=/usr/ucb/vi
set record=~/.record
set screen=20
set toplines=10
set SHELL=/bin/csh
ignore apparently-to error-to id message-id received via
```

A user can add or override mail options by creating a file named .mailrc in the home directory. A skeleton file, which contains the first three lines of the Mail.rc example shown above, is provided in /usr/skel/.mailrc.

The system administrator can copy /usr/skel/.mailrc to each new user's home directory (in its distributed form or a modified form), create the global file /usr/lib/Mail.rc, or use some combination of the two approaches.

### 2.6.4 Alias Data Base

The alias mechanism associates one or more users with a name (the alias). Mail sent to a name that is an alias is redirected to the associated users. Aliases can be nested. The mechanism is implemented on a system-wide basis through entries in the file /usr/lib/aliases, which can be edited with a standard editor. See Install_mail in Appendix B.

### 2.6.5 Encrypted Mail

Mail queued in the directory /usr/spool/mail is normally readable only by the recipient, but is in plain ASCII format. To secure mail against casual perusal, use the xget and xsend commands. Encrypted mail is queued in /usr/spool/secretmail.

## 2.6.6 Sendmail Configuration

The sendmail facility depends upon the sendmail configuration file /usr/lib/sendmail.cf for routing and other information.

Two prototype versions of this file exist in /usr/lib:

- **sendmail.main.cf** – For use on a system that routes mail in a network of systems or that connects to UUCP for mail receiving and sending.
- **sendmail.subsidiary.cf** – For use in a local system environment or when the system just receives mail from other networked systems.

Choose the template that is correct for your system.

You now must modify the template to suit your particular system configuration. Copy the proper template to a temporary file (such as /usr/lib/sendmail.custom). Edit the file and make the modifications that are indicated by comments in the file. After making the modifications, copy the file to /usr/lib/sendmail.cf.

The site-specific command script for starting Concentrix-2800 (rc.local) contains the necessary commands for starting the sendmail facility. Also, see Install_mail in Appendix B. Appendix E contains information for modifying the sendmail configuation file, and additional information concerning the operation of sendmail.

> **NOTE:** Do *not* create a frozen configuration file /usr/lib/sendmail.fc. The sendmail facility does not work correctly if a frozen configuration file exists.

## 2.7 ACCOUNTING

The system automatically collects per-process information on processor and disk use (resource accounting), and per-user information on connect time.

When sa reports the user and system times for a process executing on a complex, the base (actual) time is multiplied by the number of processors in the complex. A process which runs for 10 minutes on a complex of 6 CEs accounts for 60 minutes of CPU time.

You can customize the collection of resource accounting information, to change the built-in algorithm if you wish. See the section below, "Custom Resource Accounting" for information.

### 2.7.1 Standard Resource Accounting

The resource accounting system is controlled by the routine custom_acct.

Whenever a process exits, Concentrix-2800 calls custom_acct. Custom_acct reads process information stored in several system structures and creates an accounting record about the process. Custom_acct then writes the accounting record to an accounting file that is accumulating all of these records. Later, an application program reads the accounting file and creates a resource usage report.

Custom_acct is defined in /usr/src/sys/CONCENTRIX/custom_acct.c. The structures that custom_acct reads to find process information are defined in /usr/src/sys/h/user.h and usr/src/sys/h/proc.h. The structure of the accounting record that custom_acct writes is defined in /usr/src/sys/h/acct.h. The name of the accounting file is defined by the command accton, but is normally /usr/adm/acct.

The following line in /etc/rc turns on resource accounting and designates /usr/adm/acct as the file in which to collect the accounting information:

```
/etc/accton /usr/adm/acct
```

Change this line to turn accounting off or to designate an alternate accounting file. The following example turns accounting off:

```
/etc/accton
```

The accounting information is collected on a per-process basis and tends to make /usr/adm/acct (or the designated accounting file) grow rapidly. A line in the script /usradm/daily, which is executed once a day from /usr/lib/crontab (as distributed), executes /etc/sa (maintain system accounting files) to merge the accounting file into a summary file (/usr/adm/savacct by default). As distributed, the line is commented out. If you leave accounting on, you should uncomment this line (delete the initial pound sign); you might also want to add a line that saves previous accounting information. The following example saves one day of full accounting before the merge:

```
cp /usr/adm/acct /usr/adm/acct.yesterday
/etc/sa -s > /dev/null
```

The summary file (/usr/adm/savacct by default), of course, grows over time and should be periodically archived.

The command /etc/sa reports the information contained in /usr/adm/savacct and /usr/adm/acct, and contains a range of options. The following command produces a per-user report:

```
# /etc/sa -m > report
```

You can name an alternate accounting file on which to base the report:

```
# /etc/sa -m /usr/adm/acct.yesterday > report
```

## 2.7.2 Connect-time Accounting

The system also maintains a file containing connect times (amount of time logged in) on a per-user basis. The file is named /usr/adm/wtmp and its existence determines whether connect-time accounting is turned on. If you do not want connect-time accounting, delete this file. If you want connect-time accounting, ensure that the file exists; if you archive the file, use the touch command to create a new empty wtmp file.

Lines in the script /usr/adm/monthly, which is executed once a month from /usr/lib/crontab (as distributed), contains the following lines to archive /usr/adm/wtmp for 7 months:

```
echo "Rotating wtmp file:"
cd /usr/adm
mv wtmp.5 wtmp.6
mv wtmp.4 wtmp.5
mv wtmp.3 wtmp.4
mv wtmp.2 wtmp.3
mv wtmp.1 wtmp.2
mv wtmp.0 wtmp.1
mv wtmp   wtmp.0
cp /dev/null wtmp
```

Since the wtmp files can become fairly large, you might want to change this script to maintain fewer archives on disk.

The command /etc/ac reports connect-time information. Use the -p option to get a full report:

```
# /etc/ac -p > report
```

## 2.7.3 Custom Resource Accounting

Customizing your accounting system centers around modifying the routine custom_acct.

An experienced user can have a great deal of control over the resource accounting information that the system collects and reports. Before using custom accounting, however, there are several important points that you should keep in mind.

- You do not need to follow the instructions below in order to have resource accounting on your system. The Concentrix-2800 standard accounting package, sa, is perfectly adequate for many users. If sa is acceptable to you, you do not need to do anything else.

- The way in which you invoke custom accounting requires you to rebuild and reinstall the operating system kernel (/unix). **You must be willing to take the time to do this properly, or the entire computer system will not work correctly.**

- After you customize the accounting system, sa will not work at all.

Follow these steps to customize the accounting software:

1. Log onto the console terminal as root.

2. Edit custom_acct.c and acct.h so that custom_acct writes the kind of accounting record you want. Custom_acct may take as input any information in the user and proc structures. acct.h may define an accounting record with any structure you choose.

   (Remember that acct.h defines the output of custom_acct and therefore the structure of the accounting file.)

   (Do not change the names of any of these files, the directories that the files are stored in, or the structures in user.h and proc.h.)

3. Build and install a new Concentrix-2800 kernel (see Appendix B, Build_kernel). The new kernel will contain the new custom_acct routine.

4. Write an application program that reads the accounting records in the accounting file (usually /usr/adm/acct) and generates whatever accounting reports you want.

## 2.8 FILE SYSTEM DISK QUOTAS

The system administrator can set quotas on disk usage on a per-user, per-file system basis. File systems can be limited by the amount of space (blocks) used and by the number of files (inodes) used. The user is subject to two limits in each category: the hard limit stops any operation that would cause usage of the file system to exceed a block or inode limit; the soft limit allows the operation to continue but issues a warning message. A user exceeding a soft limit has a time limit in which to bring usage of the file system back below the soft limit, or be prevented from further write access to the file system until usage is so reduced.

The following example demonstrates quotas for four users on two file systems. All users are restricted to 10MB (hard limit) on /usr and receive a warning at 8MB (soft limit); after receiving a warning, a user has one day to reduce usage below 8MB or be refused further write access /usr. No limit is placed on the number of files than any user can have on /usr/. On / (root), user bob has the same limits as on /usr, but the other users have a hard limit of 2MB and a soft limit of 1MB; in addition, the other users have a limit of 20 files (hard limit) with a warning at 16 files (soft limit).

**Example of File System Quotas**

| | | / (root) time limit = 1 day | | /usr time limit = 1 day | |
|---|---|---|---|---|---|
| | | soft | hard | soft | hard |
| bob | blocks | 8M | 10M | 8M | 10M |
| | inodes | 0 | 0 | 0 | 0 |
| judy | blocks | 1M | 2M | 8M | 10M |
| | inodes | 16 | 20 | 0 | 0 |
| jim | blocks | 1M | 2M | 8M | 10M |
| | inodes | 16 | 20 | 0 | 0 |
| jan | blocks | 1M | 2M | 8M | 10M |
| | inodes | 16 | 20 | 0 | 0 |

A quota data base is maintained in a file named quotas in the root directory of each file system enabling quotas. Commands are supplied for editing quotas (edquota), checking and fixing quotas (quotacheck), enabling quotas (quotaon), and disabling quotas (quotaoff). In addition, the /etc/fstab file permits a quota and noquota options for automation of the enabling process. See the administration procedure Add_quotas for detailed instructions on adding file system quotas. See the specific commands mentioned above for more detail on their operation.

The administration command repquota summarizes quotas and disk usage for a specified filesystem or for all file systems using quotas. For example:

```
# repquota -v -a
/dev/xdv00r (/):
                        Block limits                        File limits
User            used    soft    hard    timeleft     used    soft    hard    timeleft
bob      +-   1080K   1024K   2048K   23.9 hours        5       0       0
jim      --    216K   1024K   2048K                     1       0       0
jan      --      0K   1024K   2048K                     0       0       0
judy     --      0K   1024K   2048K                     0       0       0
/dev/xdv00f (/usr):
                        Block limits                        File limits
User            used    soft    hard    timeleft     used    soft    hard    timeleft
bob      --   3232K   8192K 10240K                     2       0       0
jim      --    433K   8192K 10240K                     3       0       0
jan      --   1616K   8192K 10240K                     1       0       0
judy     --    432K   8192K 10240K                     2       0       0
```

The user command quota displays quotas for the specified user (only the superuser can view the limits of other users):

```
% quota -v
Disk quotas for bob (uid 116):
Filesystem     usage  quota  limit    timeleft  files  quota  limit    timeleft
/              1080K  1024K  2048K  23.9 hours      5      0      0
/usr           3232K  8192K 10240K                  2      0      0
```

# CHAPTER 3
# DEVICES

The following tables outline the commands, procedures, and files pertaining to devices.

**Administration Files**

| name | description |
|---|---|
| /dev | directory containing special files for devices |
| /etc/disktab | disk description file ** (edit) |
| /etc/fstab | static information about file systems ** (edit) |
| /etc/gettytab | terminal configuration data base ** (edit) |
| /etc/mtab | mounted file system table ** |
| /etc/stripetab | striped partition data base ** (edit) |
| /etc/ttys | terminal initialization data ** (edit) |
| /usr/src/sys/CONCENTRIX/ioconfig | input/output configuration file; source (edit) |
| /usr/src/sys/CONCENTRIX/genconf.c | program that generates autoconf.c from ioconfig |
| /usr/src/sys/CONCENTRIX/autoconf.c | input/output configuration file; generated from from ioconfig |

** See *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

**Administration Commands (Appendix A)**

| name | description |
|---|---|
| diskf | disk utility program |
| dlabel | generate disk labels from prototype file |
| extendswap | enable/disable multiple-disk swapping/paging area |
| getty | set terminal mode |
| init | system initialization and process control |
| makefs | construct a file system |
| makeramdisk | create a RAM disk image |
| mkdisknod | make special files for disks |
| mknewdevs | make default devices for initial installation |
| mknod | build special file |
| mkptynod | make special files for pseudo terminals |
| mkstripe | make striped partition |
| mktape | make special files for tapes |
| mkttynod | make special files for terminals |
| mount | mount file systems |
| tunefs | tune up an existing file system |

**Administration Procedures (Appendix B)**

| name | description |
|---|---|
| Add_terminal | add a terminal to the system |
| Extend_swap | extend area for swapping and paging |
| Format_disk | format and label disk |
| Label_disk | label (relabel) disk |

**General Commands (Concentrix-2800 Commands and Applications Manual)**

| name | description |
|---|---|
| dd | convert and copy |
| ls | list contents of directory |
| rt | read magnetic tape and print record sizes |

## 3.1 SPECIAL FILES

Special files in the directory /dev identify devices that you can access in Concentrix-2800. Writing to a special file writes to the device associated with the file and reading from a special file reads from the device associated with the file. For example, if the file /dev/rxtv00h is associated with the tape drive on unit 0 of controller 0 in 6250bpi mode, the following tar (tape archive) command writes to the device in that mode:

```
% tar cf /dev/rxtv00h *
```

The elements of a special file are as follows:

- **Major number** – Specifies the device type. The major number must be appropriate for the device type being accessed. For example, major number 0x80 (128 decimal) identifies the Xylogics VME tape controller; any special file that accesses a device of that type must have that major number. Major numbers are 8 bits, from 0x0 – 0xff (0–255 decimal). The table at the end of this section lists the major numbers for the devices supported by Concentrix-2800.

- **Minor number** – Specifies device specific information, such as the controller number, unit number, disk partition, and tape density. Minor numbers are 24 bits.

The following figure illustrates the specification of the minor number for Xylogics disk controllers:

**Example: Minor Number for Xylogics Disk Controller**

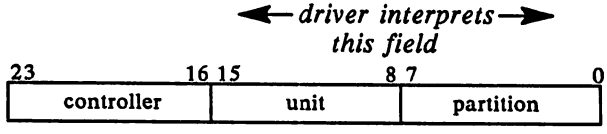


| | | | | |
|---|---|---|---|---|
| *examples:* | 0x00 | 0x00 | 0x01 | (rxdv00b) |
| | 0x00 | 0x00 | 0x16 | (rxdv00w) |
| | 0x01 | 0x01 | 0x01 | (rxdv11b) |

The following figure illustrates the specification of the minor number for Xylogics tape controllers:

**Example: Minor Number for Xylogics Tape Controller**

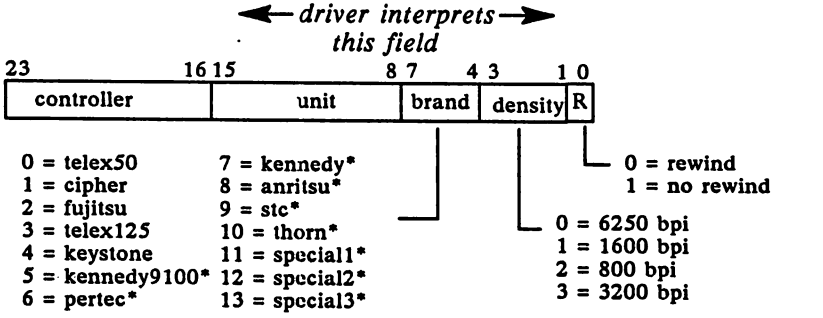


| | | | | | | |
|---|---|---|---|---|---|---|
| *examples:* | 0x00 | 0x00 | 0x1 | 0x0 | (rxtv00h) | (cipher) |
| | 0x00 | 0x00 | 0x1 | 0x1 | (nrxtv00h) | (cipher) |

The following figure illustrates the specification of the minor number for the Jaguar SCSI host adapter/tape device:

**Example: Minor Number for Jaguar SCSI Host Adapter/Exabyte 8 mm Cartridge Drive**

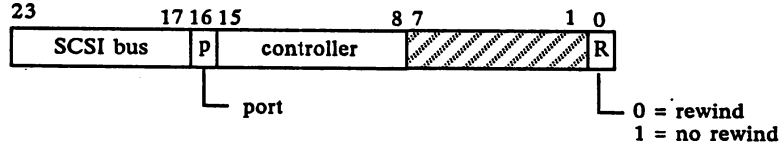


| | | | | |
|---|---|---|---|---|
| *examples:* | 0x00 | 0x02 | 0x00 | (rschs02) |
| | 0x00 | 0x02 | 0x01 | (nrschs02) |

- **Access mode** – Specifies block (also called formatted) or character (also called raw) access mode.

- **Device name** – Specifies a name for the special file. The name is arbitrary but certain conventions are highly recommended. For example, the name rxtv00h typically specifies a Xylogics VME tape device, major number 0x80 (128), controller 0, unit 0, and high density (6250bpi). The table at the end of this section lists the conventional names for the devices supported by Concentrix-2800.

You can make additional special files for one device, or make links to the special file for a device, to access the device through other names. For example, you might make the name /dev/tape0 a link to /dev/rxtv00h:

```
# ln -s /dev/rxtv00h /dev/tape0
```

### 3.1.1 Making Special Files

The mknod command makes a special file. Specify as arguments the device name, the access mode, the major number, and the minor number. The following example makes a special file named rxtv00h in character (raw) mode with a major number of 128 and a minor number of 0:

```
# /etc/mknod /dev/rxtv00h c 0x80 0
```

Concentrix-2800 also provides four commands that simplify the making of special files for certain standard devices. These commands ensure correctness and enforce naming conventions.

- The script /dev/mkdisknod makes a special file for standard and SCSI disks, given the type of controller, the controller number, and the unit number.

- The script /dev/mkptynod makes special files for pseudo terminal groups (master and slave), given the number of groups to create, or the specific group to create.

- The script /dev/mktape makes a special file for a standard tape, given the name of the kernel (usually /unix), the type of controller, the controller number, the unit number, and the type of tape drive.

- The script /dev/mkttynod makes a special file for a standard terminal, given the type of multiplexor, the multiplexor number, and the modem option.

The following example makes a special file named rxtv00h in character (raw) mode with a major number of 128 and a minor number of 0 for a cipher tape drive:

```
# /usr/alliant/mktape 0 0 cipher
```

You can examine the characteristics of a special file by listing the file in long format (ls -l). The first character of the line is c for character (raw) mode or b for block mode. The two columns preceding the date contain the major and minor numbers, respectively. The last column is the name. The following partial example lists devices whose names begin with *xdv* (typically disk devices in block mode). The minor numbers (controller, unit, and partition numbers) vary.

## 3.1.2 Listing Special Files

```
# ls -l /dev/xdv*
brw-r--r--  1 root     0x60,     0x1 Jul  4  1976 /dev/xdv00b
brw-------  1 root     0x60,     0x3 Jul  4  1976 /dev/xdv00d
brw-------  1 root     0x60,     0x5 Jul  4  1976 /dev/xdv00f
brw-------  1 root     0x60,     0xf May  8 06:54 /dev/xdv00p
brw-------  1 root     0x60,    0x10 May  8 06:54 /dev/xdv00q
brw-------  1 root     0x60,    0x11 Jul  4  1976 /dev/xdv00r
brw-------  1 root     0x60,    0x12 May  8 06:54 /dev/xdv00s
brw-------  1 root     0x60,    0x16 May  8 06:54 /dev/xdv00w
brw-------  1 root     0x60,    0x17 May  8 06:54 /dev/xdv00x
brw-------  1 root     0x60,    0x19 Jul  4  1976 /dev/xdv00z
brw-r--r--  1 root     0x60,   0x101 Jul  4  1976 /dev/xdv01b
access mode            major    minor              name
```

## 3.1.3 Concentrix-2800 Supported Devices

The following table lists the devices supported by Concentrix-2800. The major and minor numbers must be specified exactly when making a device. The name in column 1 is recommended but is not mandatory. Specify the variables as follow:

- *b* (brand) – Under minor number, specify the name brand of the tape drive: 0x0 (telex50), 0x1 (cipher), 0x2 (fujitsu), 0x3 (telex125), 0x4 (keystone).
- *c* (controller) – Under name, specify *c* as a decimal digit(s). Under minor number, specify *cc* as a hexadecimal value.
- *d* (tape density) – Under name, specify *d* as h, (6250bpi) m (1600bpi), l (800bpi), or x (3200bpi). Under minor number, specify *d* as 0x0 (6250bpi), 0x2 (1600bpi), 0x4 (800bpi), or 0x6 (3200bpi); add 1 for no rewind.
- *l* (logical unit number) – Under name, specify *l* as a single digit before the partition letter; *l* is normally 0, in which case it is omitted. Under minor number, specify as a hexadecimal value (the upper 3 bits of the lowest byte).
- *p* (disk partition) – Under name, specify *p* as a single alphabetic character. Under minor number, specify *pp* as a hexadecimal number, where a = 0x00, b = 0x01, ..., z = 0x1a.
- *ss* (SCSI bus) – Under name, specify *ss* as an integer (0 – 47). Under minor number, specify the number of the adaptor (bits 7:1) and the port on the adaptor (bit 0).
- *t* (terminal unit) – Under name and minor number, specify *t* as a hexadecimal number.
- *tt* (pseudo terminal unit) – Under name, specify the first *t* as a letter, where 0x0 = p, 0x1 = q, and so on; specify the second *t* as a hexadecimal value. Under minor number, specify *tt* as a hexadecimal value.
- *u* (unit) – Under name, specify u as a decimal digit(s). Under minor number, specify *uu* as a hexadecimal value.

**Supported Devices**

| name | mode | major | minor | description |
|---|---|---|---|---|
| xdv*cup* | block | 0x60 (96) | 0x*ccuupp* | VMEbus disk controller |
| rxdv*cup* | character | 0x60 (96) | 0x*ccuupp* | VMEbus disk controller |
| rxtv*cud* | character | 0x80 (128) | 0x*ccuubd* | VMEbus tape controller |
| nrxtv*cud* | character | 0x80 (128) | 0x*ccuubd* | VMEbus tape controller; no rewind |
| rmdk | block/chr | 0x89 (137) | 0x0 | physical memory disk |
| mlp*c* | character | 0x90 (144) | 0x*cc*0001 | VMEbus printer |
| imua*cu* | character | 0x96 (150) | 0x*uu* | image memory unit adapter card |
| tty*cu*[1] | character | 0x98 (152) | 0x*cc*000*u* | VMEbus multiplexor terminal |
| tty*cu*[1] | character | 0x98 (152) | 0x*cc*008*u* | VMEbus multiplexor term/modem |
| rschs | character | 0xa5 (165) | 0x200 | 8mm cartridge; default |
| nrschs | character | 0xa5 (165) | 0x201 | 8mm cartridge; default; no rewind |
| rschs*u* | character | 0xa5 (165) | 0x*sscc*00 | 8mm cartridge |
| nrschs*u* | character | 0xa5 (165) | 0x*sscc*01 | 8mm cartridge; no rewind |
| rsd*ssu*[*l*]*p* | character | 0xa5 (165) | 0x*ssuu*[*l*]*p* | SCSI disk device |
| sd*ssu*[*l*]*p* | block | 0xa5 (165) | 0x*ssuu*[*l*]*p* | SCSI disk device |
| console | character | 0xc1 (193) | 0x0 | port for console |
| ttyR0 | character | 0xc1 (193) | 0x1 | remote port |
| ttyC*u* | character | 0xc1 (193) | 0x*uu* | additional console port |
| ttyR*u* | character | 0xc1 (193) | 0x*uu* | additional remote port |
| tty | character | 0xc2 (194) | 0x0 | process's terminal |
| mem | character | 0xc3 (195) | 0x0 | main memory |
| kmem | character | 0xc3 (195) | 0x1 | kernel memory |
| null | character | 0xc3 (195) | 0x2 | null device |
| full | character | 0xc3 (195) | 0x6 | full device |
| tty*tt* | character | 0xc4 (196) | 0x*tt* | pseudo terminal; server |
| pty*tt* | character | 0xc5 (197) | 0x*tt* | pseudo terminal; client |
| stripe*u* | block | 0xc8 (200) | 0x*uu* | striped device |
| rstripe*u* | character | 0xc8 (200) | 0x*uu* | striped device |
| realtime | character | 0xcc (204) | 0x0 | realtime protection device |
| ex*c* | none[2] | 0xcd (205) | 0x*cc* | Excelan Ethernet controller |

[1] Name terminal 0x*cc*000*u* tty*cu*.nomodem or name terminal 0x*cc*008*u* tty*cu*.modem. The unextended name appears in ttys.

[2] Access to network drivers is via kernel routine algorithms; these devices do not have entries in /dev.

(continued)

**Supported Devices (continued)**

| | | | | |
|---|---|---|---|---|
| enc | none[2] | 0xce (206) | 0x*cc* | CMC-130 Ethernet controller |
| klog | block/chr | 0xcf (207) | 0x0 | log device for syslogd |
| uvm*c* | block/chr | 0xd0 (208) | 0x*cc* | Ultranet hardware driver |
| unet*c* | block/chr | 0xd1 (209) | 0x*uu* | Ultranet protocol driver |
| unet | none[2] | 0xd1 (209) | 0 | Ultranet internet driver |
| npd | character | 0xd2 (210) | 0x0 | NFS remote execution driver |
| pg*c* | none[2] | Oxd3 (211) | Ox*cc* | Interphase Peregrine FDDI |
| pg*c* | character | Oxd4 (212) | Ox*cc* | Peregrine station management |
| ge | character | 0xda (218) | 0x0 | DNX-2800 DECnet driver |
| bpf*u* | character | oxdb (219) | Ox*uu* | Berkeley Packet Filter |
| imb*u* | character | 0xe0 (224) | 0x*uu* | byte interface to Diagnostix-2800 |

[2] Access to network drivers is via kernel routine algorithms; these devices do not have entries in /dev.

## 3.1.4 Input/output Address Ranges

The following table specifies the start of the input/output address range (in hexadecimal) allotted to devices of each type, the increment between addresses (in hexadecimal) in the range, the number of addresses (in decimal) in the range, and the address modifiers (for programmed input/output, direct memory access (DMA via fixed maps), and DMA via dynamic maps).

**Input/output Address Ranges**

| base | increment | number | modifiers* | device type |
|---|---|---|---|---|
| na | na | 1 | VME_PRIV_DATA_AM<br>VME_PRIV_DATA_EXT_AM | imu - IMUAC |
| 0x4000 | 0x800 | 4 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | ve - DMA exerciser |
| 0x7740 | 0x20 | 4 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | uvm - Ultranet |
| 0xc000 | -0x800 | 12 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | scsi - SCSI |
| 0xe400 | 0x200 | 4 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | mvp/mlp printer/plotter |

(continued)

**Input/output Address Ranges (continued)**

| | | | | |
|---|---|---|---|---|
| 0xec00 | 0x10 | 16 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | xtv - tape |
| 0xed00 | 0x10 | 8 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | xmv - multiplexor |
| **base** | **increment** | **number** | **modifiers*** | **device type** |
| 0xee40 | 0x10 | 32 | VME_PRIV_SHORT_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | xdv - disk |
| 0xe00000 | 0x40000 | 6 | VME_PRIV_DATA_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | en - CMC Ethernet |
| 0xf90000<br>(0xf80000 | 0x8000<br>for Diagnostix-2800) | 4 | VME_PRIV_DATA_AM<br>VME_PRIV_DATA_EXT_AM<br>VME_PRIV_DATA_EXT_AM | ex - Excelan Ethernet |

* See /usr/include/machine/vme.h

## 3.1.5 Diagnostix-2800 Supported Devices

The following table lists the devices supported by Diagnostix-2800. The major and minor numbers must be specified exactly when making a device. The name in column 1 is recommended but is not mandatory. Specify the variables as follow:

- *b* (port on SCSI adaptor) Use *b* as an integer value in calculating the minor number (bit 7).
- *c* (controller) – Under name, specify *c* as a decimal digit(s). Use *c* as an integer value in calculating the major number.
- *u* (unit) – Under name, specify *u* as a decimal digit(s). Use *u* as an integer value in calculating the minor number.
- *p* (disk partition) – Under name, specify *p* as a single alphabetic character. Use *p* as an integer value in calculating the minor number, where a=0, b=1, ..., z=25.
- *d* (tape density) – Under name, specify *d* as h (6250bpi), m (1600bpi), l (800bpi), or x (3200bpi). Use *d* as an integer value in calculating the minor number, where h=0, m=32, l=64, and x=96.
- *l* (logical unit number) – (SCSI devices) Under minor number, specify as either 0 or 1 (bit 3).
- *ss* (SCSI bus) – Under name, specify *ss* as an integer (0 – 6). Under minor number, specify the number of the SCSI controller (bits 4:6).

  **NOTE:** SCSI minors 0 – 15 and 128 – 128+15 do not conform to the normal SCSI minor layout. Each number has a unique meaning (the Exabyte helical scan 8mm cartidge tape drive, rewind and no rewind versions): rschs and nrschs.

**Diagnostix-2800 Supported Devices**

| name | mode | major | minor | description |
|---|---|---|---|---|
| xdv*cup* | block | 96+*c* | 32*u*+*p* | VMEbus disk controller |
| rxdv*cup* | character | 96+*c* | 32*u*+*p* | VMEbus disk controller |
| rmdk | block/chr | 137 | 0 | physical memory disk |
| rschs | character | 165+*c* | 0 | 8mm cartridge; default |
| nrschs | character | 165+*c* | *bsslp* | 8mm cartridge; default; no rewind |
| rschs*u* | character | 165+*c* | *bsslp* | 8mm cartridge |
| nrschs*u* | character | 165+*c* | *bsslp* | 8mm cartridge; no rewind |
| rsd*ssup* | character | 165+*c* | *bsslp* | SCSI disk device |
| sds*sup* | block | 165+*c* | *bsslp* | SCSI disk device |
| console | character | 193 | 0 | duart port for console |
| ttyR0 | character | 193 | 1 | duart port for modem (remote port) |
| tty | character | 194 | 0 | process's terminal |
| mem | character | 195 | 0 | main memory |
| kmem | character | 195 | 1 | kernel memory |
| null | character | 195 | 2 | null device |
| full | character | 195 | 6 | full device |
| log | block/char | 207 | n/a | log device for syslogd |
| imb | character | 224, | 0 | imb0 Concentrix-2800 console port |
| imb | character | 224 | 1 | imb1 background Diagnostix-2800 shell; processes requests from dsh |
| imb | character | 2244 | 2,3 | unused, available for Diagnostix-2800 & Concentrix-2800 communication |

On Diagnostix-2800, the SCSI adaptor number is part of the device major (each successive SCSI adaptor gets the next higher major number), and the 8 bits of the SCSI minor are allocated as follows:

**Minor Number for Diagnostix-2800 SCSI device**

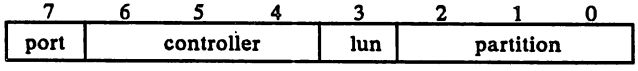


| | | |
|---|---|---|
| *examples:* | 165, 112 | rsd00b |
| | 166, 112 | rsd20b |

Notice that Diagnostix-2800 is restricted to accessing only two logical units (0 and 1), and only eight partitions. The eight possible partitions are as follows:

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| partition number: | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| partition letter: | b | d | f | h | p | q | r | z |

SCSI id zero is encoded as a 7 in order to remain compatible with the use of the small minor numbers for accessing the Exabyte tape drive.

The command mkdisknod makes the appropriate special files.

## 3.1.6 Default Special Files

Concentrix-2800 provides the following script for automatically making, at installation time, the special files to support a basic device configuration:

```
/dev/mknewdevs
```

If your device configuration differs from the basic device configuration, you can modify the script. The script contains comments indicating how modifications can be made.

The command syscfg -d displays the devices attached to each IOM on your system. The information provided includes the device name, the major number, the input/output address, and a description.

## 3.2 DISK DEVICES

Alliant currently supports the following disk controllers:

- **Xylogics 754 (VMEbus)** disk controller – Supports the SMD-E interface at a transfer rate of 3.0MB/second.
- **Xylogics 7800 (VMEbus)** disk controller – Supports the IPI-2 interface at a transfer rate up to 10.0MB/second.

### 3.2.1 Disk Names

The conventional names for disk devices (special files) are as follows:

**/dev/xdv***cup* VMEbus, block (formatted)

**/dev/rxdv***cup* VMEbus, character (raw) mode

Where:

*c* is an integer 0 – 11 that specifies a disk controller.

*u* is an integer 0 – 7 that specifies a unit (drive).

*p* is a single alphabetic character that specifies a partition.

### 3.2.2 Supported Disk Drives

The following disk drives are supported by Concentrix-2800 as shipped:

- **fj690s** – (SMD) Fujitsu M2344 8-inch 690MB disk drive.
- **fj1000s** – (SMD) Fujitsu M2382 8-inch 1GB disk drive.
- **fj2600s** – (IPI-2) Fujitsu M2671P 8-inch 2.65GB disk drive.

## 3.2.3 Disk Geometry

The following terms are presented for review:

**Terminology**

| | |
|---|---|
| platter | Most disk devices consist of several platters mounted on a spindle spinning at high speed. Data is stored on a thin magnetic coating. |
| head | To read and write information on the platters, several heads are mounted on a single arm so that they travel together. Usually, there are two heads for each platter (one for the top surface and one for the bottom). The arm moves over the surface of the disk. This movement is called seeking. The positions of the heads are controlled by hardware. |
| track | The portion of a disk passing under a single, stationary head during a single rotation is called a track. |
| cylinder | The set of tracks for all the heads at a single seek position is called a cylinder. |
| sector | Each track is divided into segments called sectors. To read successive sectors in the same track, the heads remain stationary as the disk spins under them. The disk drive generates a timing signal that identifies the sectors for the disk controller. A sector can be 512, 1024, 2048, or 4096 bytes. |
| controller | A disk controller is a board that plugs into a bus and communicates directly with one or more disk drives. It handles error checking, data transfer, and the arrangement of the data on the disk. A controller and its associated devices are identified to the system at configuration time and an appropriate driver module is compiled into the system. When the system boots, the autoconfiguration facilities probe for the controller on the bus and, if found, enable the software support for it. |
| unit | A disk controller can have more than one drive attached to it. The individual disk drives are called units. |
| driver | A disk driver is a software module that interfaces the operating system to a specific type of controller. It arranges for buffers, handles interrupts from the controller, handles disk errors, keeps track of I/O requests from programs, and takes care of other disk functions. |

### 3.2.3.1 Disk Storage Layout

Alliant forms non-SCSI disks with a fixed (constant) number of sectors per track, and a fixed number of sectors per cylinder. The following algorithm applies to non-SCSI disks:

```
sectors_per_cylinder = sectors_per_track * tracks_per_cylinder
```

The storage model for a SCSI disk is simply a linear array of blocks (of a user specified size) numbered 0, 1, 2, ..., *end*. Although the disk contains platters divided into cylinders, heads, and sectors, this information is, for the most part, hidden from the operating system.

No standard way exists to lay out the numbered blocks on the physical media of the disk. SCSI disk manufacturers are free to allocate varying numbers of sectors per track or per cylinder, to handle disk defects, increase capacity, and improve performance. Fujitsu, for instance, allocates a fixed number of sectors per track except for the last track of each cylinder, which has fewer sectors in order to provide room for spares. (The spares are used for bad spots within that cylinder). The number of sectors per cylinder is constant.

One vendor (attempting to increase storage capacity) allocates more sectors per track on the outer cylinders of the disk since the tracks there are larger. You cannot calculate which cylinder and head a particular sector lies on without a table lookup based on the actual sector number.

Since the SCSI specification for flaw maps provides a nearly standard way of dealing with media defects, Alliant has opted to use a subset of these features which appears to be supported by all drives.

Normally Alliant locates SCSI flaws in "bytes from index" format. This is the format used by SMD and IPI-2 drives; the flaw location is expressed as a cylinder, head, and number of bytes from the start of the track (the "index" pulse). Certain drives only support "physical sector" format in which the flaw is expressed as a cylinder, head, and sector number. In this case, an entry must be made in the scsi_ids array to recognize the particular drive and indicate that the driver is to use this alternate format. Alliant does not support "logical sector" format.

### 3.2.3.2 Disk Labels

Under Concentrix-2800, each physical disk contains a label at a standard location (sector 0 of the boot partition). The label describes the physical characteristics of the disk. The administration command /etc/dlabel places a label on a disk through reference to the prototype labels in the disk description file /etc/disktab. To examine the prototype labels in the file /etc/disktab, type **more /etc/disktab**. To examine the label on a disk, apply the S command of /etc/diskf to the z partition of the disk:

```
# /etc/diskf /dev/rxdv01z S
```

Alliant supplies a disk description file that contains a standard description for each supported disk. A standard label is placed on a disk as part of the formatting process; see Format_disk in Appendix B. Typically, you do not perform this procedure because Alliant formats and labels disks before shipping. You can add an entry to the disk description file and apply the entry to a disk, for example, to support non-standard partitioning of disks; see Label_disk in Appendix B. You should not delete entries from the disk description file.

### 3.2.3.3 Partitions

Under Concentrix-2800, each physical disk is divided into logical partitions (up to 128); each partition is named for a letter of the alphabet. Each partition is referenced in each mode (block and character) by a distinct special file in the /dev directory.

Two root and swap partitions accommodate the two different operating systems: Concentrix-2800 (which executes on the FX/2800 processors) and Diagnostix-2800 (which executes on the VMEDP).

The following table lists the partitions defined in the standard disk description file.

**Standard Partitions**

| | |
|---|---|
| b | boot partition. Do not change this partition in any way unless specifically instructed to do so by Alliant Customer Services. |
| r | typically used as the root partition for Concentrix-2800. The root partition is reserved for system software, including the Concentrix-2800 kernel, and must be accessible to the master IOM (IP0) and to the master VME-IP on IOM0. |
| s | typically used as a minimal swapping/paging partition for Concentrix-2800. |
| f | typically used as a minimal user file partition. |
| q | typically used as the swapping/paging partition for Diagnostix-2800. |
| p | typically used as the root partition for Diagnostix-2800. |
| d | diagnostic partition. Do not change this partition in any way unless specifically instructed to do so by Alliant customer services. |
| h | typically used as a large swapping/paging or user file partition on disks that do not contain a root file system. Overlaps the r, s, f, q, and p partitions. |
| w | typically used to access the entire disk when making a copy of it or when using the diskf and dlabel commands. Overlaps all partitions except d and t. **Note:** Now obsolete; use z. |
| z | typically used to access the entire disk when making a copy of it or when using the diskf and dlabel commands. Overlaps all partitions except t. |
| t | cylinder for alternate tracks (1 and 2GB non-SCSI drives only). Do not access this partition directly. The diskf alternate track commands access it internally. |

The table below shows the approximate size in megabytes of the standard partitions on the supported disks:

**Partition Sizes**

| type | b | d | r | s | f | q | p | h | w | z | t |
|---|---|---|---|---|---|---|---|---|---|---|---|
| fj690s | 1 | 1 | 38 | 173 | 244 | 10 | 57 | 524 | 525 | 526 | na |
| fj1000s | 1 | 1 | 40 | 299 | 370 | 10 | 57 | 779 | 780 | 781 | 1 |
| fj2600s | 1 | 1 | 43 | 876 | 876 | 11 | 60 | 1966 | 1967 | 1968 | 1 |

The following diagram shows how these partitions are typically laid out on an fj690s disk:

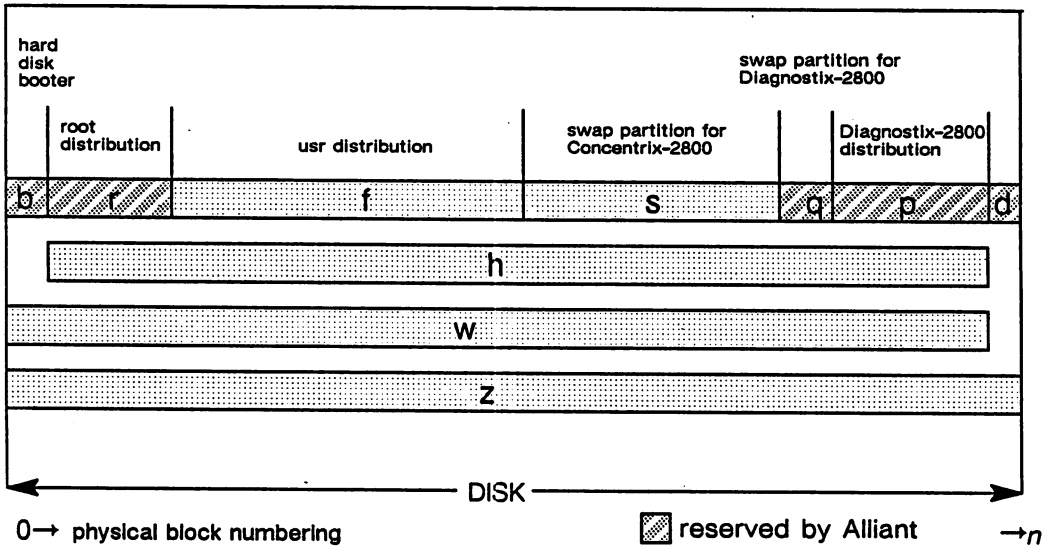


The following diagram shows how these partitions are typically laid out on an fj1000s or fj2600s root disk:

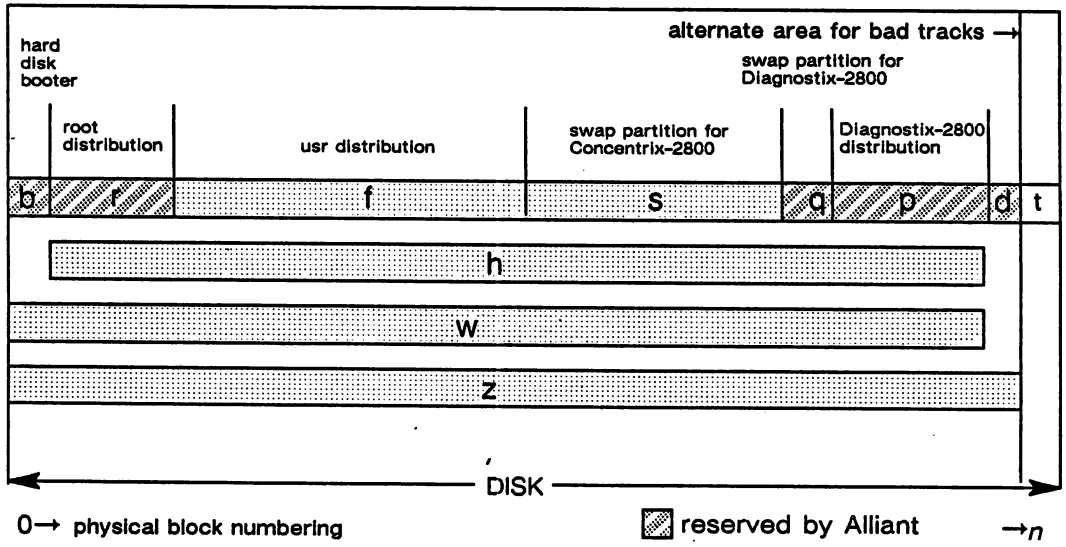

### 3.2.3.4 Flaws

A disk as delivered typically has a number of bad sectors; the disk manufacturer notes the bad sectors on a *hardware flaw map* on the disk. Prior to formatting a disk, the hardware flaw map is moved to a *flaw file* (an ordinary ascii file) with the p command of /etc/diskf. After formatting the disk, the flaw file is written to a special location on the disk called the *internal flaw map* with the D command of /etc/diskf. Typically, you save the flaw file in case the internal flaw map is corrupted; typically, you name the flaw file after the disk type and serial number of the drive, joined by an underscore. For example, the flaw file for a Fujitsu M2344, serial number 12345, is typically named fj690s_12345.

If you format a new disk (not typical), follow the sequence described in Format_disk in Appendix B exactly. You must format the disk, label the disk, and move the flaw information in proper sequence.

During the formatting process, the flaw information is used to substitute spare sectors for the bad sectors. Each track contains two (1GB disks) or four (fj690s) spare sectors at the end; if a bad sector exists, the bad sector and the remaining sectors on the track are logically slipped towards the end of the track.

If a sector on an SMD or IPI-2 disk goes bad after formatting, you can manually slip the bad sector. A bad sector is typically noticed through receipt of a "hard ecc error" message. The message appears on the terminal of the user accessing the disk at the time and provides the device name and other information. Part of the information provided is the beginning location of the operation that encountered the error in the form:

```
<155.,6.,16.>
```

Where 155 is the cylinder, 6 is the head, and 16 is the sector. This number is sometimes inaccurate; to determine the bad sector precisely, always read the cylinder and head with the r command of /etc/diskf:

```
# /etc/diskf /dev/rxdv01z r 155.6 > /dev/null
```

The diskf utility notes the bad sector.

Then slip the sector with the s command of /etc/diskf:

```
# /etc/diskf /dev/rxdv01z s 155.6.16
```

If the s command fails because the spare sectors on the track are already used, map the entire track out to the t partition (fj1000s disks only):

```
# /etc/diskf /dev/rxdv01z am 155.6
```

Update the flaw information with the following sequence of commands:

1. Copy the internal flaw map to a flaw file:

   ```
   # /etc/diskf /dev/rxdv01z B > fj690s_12345
   ```

2. Make a file of slipped sectors in flaw file format:

   ```
   # /etc/diskf /dev/rxdv01z b 155.6 > flaws.tmp
   ```

3. Append the slipped sector information to the flaw file:

```
# /bin/cat flaws.tmp >> fj690s_12345
```

4. Write the slipped sector information to the internal flaw map:

```
# /etc/diskf /dev/rxdv01z D fj890s_12345
```

5. Remove the temporary file:

```
# rm flaws.tmp
```

### 3.2.4 Disk Layout Considerations

Layout considerations involve three major categories of disk space:

- **root** – The root file system contains the core system files and is the parent of all other file systems. As shipped by Alliant, the r partition on the first disk (controller 0, unit 0) contains the root file system for Concentrix-2800; the raw device name is /dev/rxdv00r. The p partition of the first disk contains the root file system for Diagnostix-2800; the raw device name is /dev/rxdv00p.

  The installation tape only supports access to the following four root disks: /dev/rxdv00, /dev/rxdv01, /dev/rxdv10, and /dev/rxdv11. Once the system is booted off an installed disk, all controllers and units may be accessed.

  You can modify the locations of the root partitions and add alternate root partitions by using the dlabel -Cr or -Dr command. Root partitions must be on devices connected to IP0. Log the location(s) of the root partition(s) if you make modifications.

- **swap** – The swapping/paging area provides space for physical memory overflow. As shipped by Alliant, the s partition on the first disk (controller 0, unit 0) contains the swapping/paging area for Concentrix-2800; the raw device name is /dev/rxdv00s. The q partition of the first disk contains the swapping/paging area for Diagnostix-2800; the raw device name is /dev/rxdv00q. Sometimes this space is not sufficient; see the guidelines below.

  You can modify the locations of the swap partitions and add alternate swap partitions by using the dlabel -Cs or -Ds command. Log the location(s) of the swap partition(s) if you make modifications.

- **user** – The remaining disk space (except for the boot and diagnostic partitions) is used for user file systems. Typically, the f partition on the first disk (controller 0, unit 0) and the h partitions on other disks contain the user file systems; the raw device names are /dev/rxdv00f, /dev/rxdv01h, /dev/rxdv10h, and so on. However, additional swapping/paging requirements may require modifications to this scheme; see the guidelines for swapping/paging below. If the sizes of the standard partitions do not fit your file system needs, you may want to relabel a disk or disks; see the guidelines below.

### 3.2.4.1 Protecting Disks

Ensure that devices, that is, the special files in /dev, are protected properly. The protection code should be octal 600, read and write access for root only.

Exercise extreme caution in accessing disk devices as superuser. When applying commands such as diskf, dlabel, and makefs, be certain that the device name is correct – not only the controller and unit numbers, but the partition. Writing to or initializing a partition that overlaps another partition destroys the second partition. Executing /etc/makefs on /dev/rxdv00z, for example, would destroy /dev/rxdv00r and /dev/rxdv00f.

Removing unused partitions, as explained in the next section, reduces the risk of accidentally destroying data by writing to an overlapping partition. However, the w partition cannot be removed.

A safety measure you might employ when working on a disk partition is to remove the protection on the special file (change the protection to octal 666), then access the device as an ordinary user, not superuser. In this way, you cannot accidentally access another device. Be sure to change the protection back to octal 600.

Another safety measure you can employ is to set the environment variable DISKF_SCRATCH_DISK and let diskf default to this, instead of specifying the disk explicitly. The following example sets the diskf default to /dev/rxdv01z, then that disk is reinitialized.

```
# setenv DISKF_SCRATCH_DISK /dev/rxdv01z
# diskf B >fj1000s_99999
# diskf f fj1000s_99999
# diskf D fj1000s_99999
# diskf C  internal
```

### 3.2.4.2 Unmarking Partitions and Removing Unused Partitions

Disks shipped by Alliant typically have the r partition marked as the root partition and the s partition marked as a swap partition. If you boot the system and the disk that contains the actual swap partition is offline, a partition on another disk might accidentally be selected as the swap partition, which would destroy any data in any overlapping partitions. Therefore, you should relabel the disks that do not contain the root and swap partitions, omitting the -Cr, -Dr, -Cs, and -Ds options. For example, if you have three Fujitsu M2344 disks mounted on controller 0, units 0 and 1, and controller 1, unit 0, and the root and swap partitions are on the first disk (as shipped), you would issue the following commands (using standard names):

```
# /etc/dlabel -iw fj690s /dev/rxdv01z
# /etc/dlabel -iw fj690s /dev/rxdv10z
```

Removing unused partitions lowers (but does not eliminate) the chance of accidentally destroying data by writing to an overlapping partition. You cannot remove the following partitions: b, d, t, and z (which encompasses the entire disk except for the t partition). See Label_disk in Appendix B.

### 3.2.4.3 Swapping and Paging Considerations

The standard configuration, which uses the s partition on the root disk for swapping and paging, usually suffices. However, a large number of processes or processes with large virtual memory requirements may overload the standard swapping/paging partition. If the paging area begins to fill up, messages to that effect appear on the console. If the swap area overflows, a message to that effect appears on the console after the fact.

You might also write the following or a similar shell script to periodically display the usage of the swap partition:

```
#!/bin/csh -f
#
# Script to determine the amount of page and swap disk space
# currently in use.
#
adb /unix /dev/kmem <<EOF > cwfg-temp
primary+0x4/D
primary+0x14/D
swdevt/D
swdevt+0x4/D
EOF
set pars = `cat cwfg-temp`
/bin/rm -f cwfg-temp
@ i = $pars[4] * 100
@ i = $i / $pars[2]
@ j = $pars[8] * 100
@ j = $j / $pars[6]
@ k = 100 - $i
@ l = 100 - $j
echo ""
echo "     Disk Page/Swap File Space Currently In Use:"
echo "          page area = $k% full"
echo "          swap area = $l% full"
echo ""
```

The following example invokes the script (its name is swpg):

```
# swpg
     Disk Page/Swap File Space Currently In Use:
        page area = 0% full
        swap area = 40% full
```

If the use of the paging/swapping partition is high (75% or more), you should implement an alternate scheme for partitioning your disk space. (If you notice that the use of paging and swapping space is out of balance, and you are short on one or the other, you might adjust the paging and swapping system parameter swpgratio; see Patch_kernel in Appendix B. Consulting with Alliant Customer Service before making this adjustment is recommended.)

The following figure illustrates the standard configuration for the root, swap, and user partitions, and some typical alternatives. The figure assumes a system with four Fujitsu M2344 disks mounted on controllers 0 and 1, units 0 and 1, with standard names. The commands on the right side of each configuration demonstrate the commands that effect the configuration. See Extend_swap in Appendix B for the exact procedure.

**Partition Configurations**

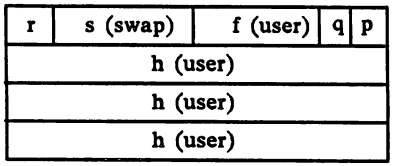


**standard configuration**

```
# /etc/dlabel -iw -Crr -Css -Drp -Dsq \
  fj690s /dev/rxdv00w
# /etc/dlabel -iw fj690s /dev/rxdv01w
# /etc/dlabel -iw fj690s /dev/rxdv10w
# /etc/dlabel -iw fj690s /dev/rxdv11w
```

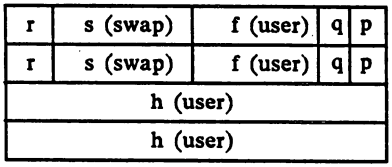


**swap on two s partitions**

```
# /etc/dlabel -iw -Crr -Css -Drp -Dsq \
  fj690s /dev/rxdv00w
# /etc/dlabel -iw -Css fj690s /dev/rxdv01w
# /etc/dlabel -iw fj690s /dev/rxdv10w
# /etc/dlabel -iw fj690s /dev/rxdv11w
```

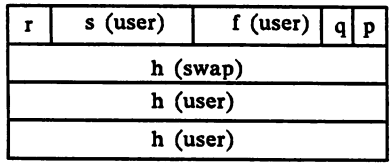


**swap on a full disk**

```
# /etc/dlabel -iw -Crr-Drp -Dsq fj690s \
  /dev/rxdv00w
# /etc/dlabel -iw -Csh fj690s /dev/rxdv01w
# /etc/dlabel -iw fj690s /dev/rxdv10w
# /etc/dlabel -iw fj690s /dev/rxdv11w
```

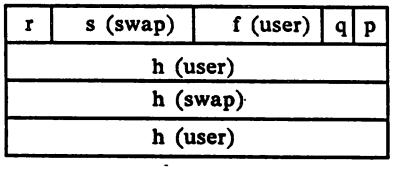


**swap on s partition plus a full disk**

```
# /etc/dlabel -iw -Crr -Css -Drp -Dsq \
  fj690s /dev/rxdv00w
# /etc/dlabel -iw fj690s /dev/rxdv01w
# /etc/dlabel -iw -Csh fj690s /dev/rxdv10w
# /etc/dlabel -iw fj690s /dev/rxdv11w
# /usr/alliant/extendswap
```

If you use more than one disk for swapping and paging, the following rules must be observed:

- One partition per disk can be marked for swapping and paging.
- At least one swap partition must be accessible to IP0.
- All disks and controllers must be of the same type.

Multiple swap partitions permit a paging/swapping area that exceeds the size of one disk, and provides some performance benefit if the partitions are spread over disks on different controllers.

### 3.2.4.4 Non-Standard User Partitions

The standard partitions provide space for a small file system on the root disk and a large file system on each of the remaining disks. (You may, of course, choose to take up some of this space to extend the swapping/paging area.) If this partitioning scheme does not fit your needs, you can relabel disks as required. For example, you might want to be able to place two file systems of equal size on one disk. See Label_disk in Appendix B, "Adding User Partitions." Also see Chapter 4, "File Systems."

## 3.3 SMALL COMPUTER SYSTEM INTERFACE (SCSI) DEVICES

Alliant currently supports the **Interphase V/SCSI 4210 Jaguar** SCSI interface, a VMEbus SCSI host adaptor. Each adaptor supports up to two SCSI busses. Each bus supports up to six SCSI controllers. (A SCSI "controller" may be a controller, in the traditional sense, that connects one or more separate devices; in this case, a logical unit number (lun) identifies each device. More typically, a SCSI "controller" is a single unit that combines a controller and a device; in this case, the logical unit number is 0.)

### 3.3.1 SCSI Disk Devices

(See also the discussion of disk geometry in the preceding section.)

The conventional names for SCSI disk devices (special files) are as follows:

| | |
|---|---|
| /dev/sd*ssu*[*l*]*p* | SCSI device, block (formatted) |
| /dev/rsd*ssu*[*l*]*p* | SCSI device, character (raw) mode |

Where:

- *ss* is an integer 0 – 47 that specifies a SCSI bus. The low order bit specifies the port and the high order bits specify the adaptor, so that bus 0 and 1 belong to host adaptor 0, bus 2 and 3 belong to host adaptor 1, and so on.
- *u* is an integer 0 – 7 that specifies a SCSI controller.
- *l* is a single digit that specifies the logical unit number of the device on the controller. This digit is omitted if it is 0 (normally the case).
- *p* is a single alphabetic character that specifies a partition.

The following figure illustrates the specification of the minor number for a SCSI disk:

**Minor Number for SCSI Disk Device**

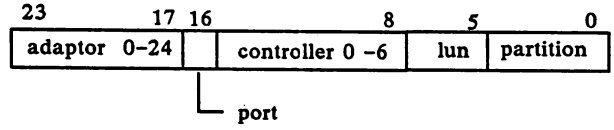


| *examples:* | | | | |
|---|---|---|---|---|
| | 0x00 | 0x00 | 0x19 | (rsd00z) |
| | 0x02 | 0x03 | 0x1 | (rsd23b) |

### 3.3.1.1 Supported SCSI Disk Drives

Alliant currently supports the Fujitsu M2266 5.25-inch 1.2GB SCSI disk drive.

The Concentrix-2800 default configuration includes the following devices (in /dev) for accessing SCSI disks:

| name | mode | major | minor | description |
|---|---|---|---|---|
| rsd00b | character | 0xa5 (165) | 0x1 | SCSI disk |
| rsd00d | character | 0xa5 (165) | 0x3 | SCSI disk |
| rsd00f | character | 0xa5 (165) | 0x5 | SCSI disk |
| rsd00h | character | 0xa5 (165) | 0x7 | SCSI disk |
| rsd00p | character | 0xa5 (165) | 0xf | SCSI disk |
| rsd00q | character | 0xa5 (165) | 0x10 | SCSI disk |
| rsd00r | character | 0xa5 (165) | 0x11 | SCSI disk |
| rsd00s | character | 0xa5 (165) | 0x12 | SCSI disk |
| rsd00w | character | 0xa5 (165) | 0x16 | SCSI disk |
| rsd00x | character | 0xa5 (165) | 0x17 | SCSI disk |
| rsd00z | character | 0xa5 (165) | 0x19 | SCSI disk |

### 3.3.1.2 SCSI Disk Labels

Under Concentrix-2800, each physical disk contains a label at a standard location (sector 0 of the boot partition). The label describes the physical characteristics of the disk. The administration command /etc/dlabel places a label on a disk through reference to the prototype labels in the disk description file /etc/disktab. To examine the label on a SCSI disk, apply the Sv command of /etc/diskf to the z partition of the disk:

```
# /etc/diskf /dev/rsd04z Sv
sample:\
        :ds=sample:\
        :bs=4096:ht=1:\
        :nc=1658:nh=15:ns=12:ss=3:as=4096:rm=3600::sc=177:\
    :0b=       0m:   :Ub=     708k: (   0.00m-   0.69m,     0.00c-   1.00c) \
    :0d=1172448k:   :Ud=1173156k: (1144.97m-1145.66m, 1656.00c-1657.00c) \
    :0f= 337716k:   :Uf=1100940k: ( 329.80m-1075.14m,  477.00c-1555.00c) \
    :0h=     708k:  :Uh=1172448k: (   0.69m-1144.97m,    1.00c-1656.00c) \
    :0p=1100940k:   :Up=1161828k: (1075.14m-1134.60m, 1555.00c-1641.00c) \
    :0q=1161828k:   :Uq=1172448k: (1134.60m-1144.97m, 1641.00c-1656.00c) \
    :0r=     708k:  :Ur=  44604k: (   0.69m-  43.56m,    1.00c-  63.00c) \
    :0s=   44604k:  :Us= 337716k: (  43.56m- 329.80m,   63.00c- 477.00c) \
    :0w=       0m:  :Uw=1172448k: (   0.00m-1144.97m,    0.00c-1656.00c) \
    :0z=       0m:  :Uz=1172448k: (   0.00m-1144.97m,    0.00c-1656.00c)
```

The procedures Format_disk and Label_disk in Appendix B include instructions for formatting and labeling SCSI disks.

### 3.3.1.3 SCSI Flaws

A disk as delivered typically has a number of bad sectors; the disk manufacturer notes the bad sectors on a *hardware flaw map* on the disk. On SCSI disks, this is called the PLIST (primary defect list); it records information by cylinder, head, and sector. When you format a SCSI disk, the drive reads its own PLIST, marks the bad spots, and formats around them. To see the PLIST ("native flaws"), use the diskf pp command as shown in the example below:

```
# diskf /dev/rsd00w pp

84      0       31049           0
453     0       32894           0
737     0       11512           0
835     0       888             0
1237    0       27045           0
        ...
```

If a sector on a SCSI disk goes bad as the drive is running, the drive recognizes the "grown" flaw and reports the block number(s) to the user in a console message. To fix the flaw, slip the sector with the s command of /etc/diskf, supplying the block number reported in the console message:

```
# /etc/diskf  /dev/rsd00w s block-number
```

The diskf utility then uses the "reassign blocks" command which fixes the flaw and records the bad sector in the drive's GLIST (the list of all "grown" flaws).

To print out the GLIST, use the diskf p command:

```
# /etc/diskf /dev/rsd00w p

0       0       7347            0
0       2       22045           0
```

You can redirect this output to a file and use that file for the diskf f format command. The drive will slip bad blocks on *both* the PLIST and the supplied list during the formatting process, and save the supplied list on the disk as the initial GLIST. Note that the diskf f command only formats *entire* SCSI disks:

```
# /etc/diskf /dev/rsd00w  p > flaw-file  [save exisitng GLIST]
# /etc/diskf /dev/rsd00w  f  flaw-file   [format with supplied list]
```

## 3.3.2 SCSI Tape Devices

Alliant currently supports the **Exabyte Model 8200 Cartridge Tape System.** This device is an 8mm Helical-Scan Tape Drive and Subsystem. Names for the 8mm cartridge devices are defined as shown below.

The conventional names for 8mm cartridge devices (special files) are as follows:

| | |
|---|---|
| **/dev/rschs***bc* | SCSI helical scan 8mm cartridge, rewind |
| **/dev/nrschs***bc* | SCSI helical scan 8mm cartridge, no rewind |

Each device requires one special file each for rewind and no-rewind modes; only character mode is allowed.

The Concentrix-2800 default configuration includes the following devices (in /dev) for accessing cartridges:

| name | mode | major | minor | description |
|---|---|---|---|---|
| rschs02 | character | 0xa5 (165) | 0x200 | 8mm cartridge, rewind |
| nrschs02 | character | 0xa5 (165) | 0x201 | 8mm cartridge, no rewind |

## 3.3.3 SCSI Switch Configuration

The Interphase VMEbus SCSI host adaptor switch configuration for /dev/rschs is as follows:

- **Bit 9** – Always off.
- **Bits 8:4** – Most significant 5 bits of the device address complemented (a switch setting of off means a bit setting of 1).
- **Bits 3:1** – The SCSI controller id complemented.

An example follows:

```
switch      9   8   7  6   5   4    3   2   1
           off off off on  on  on   off off off
           --- ----------------    -----------
                 address C000        scsi id 7
```

## 3.3.4 Loading an 8mm Cartridge

To load an Exabyte 8mm cartridge:

1. Open the door of the tape drive by pressing the LOAD/UNLOAD button on the lower right side of the drive.
2. If you want write protection, move the write-protect switch on the cartridge to the left, so that the red tab in the recessed area is visible.
3. Insert the cartridge with the write-protect switch facing you, toward your right.
4. Close the door.

The figure below shows an Exabyte cartridge drive with a tape loaded.

**Exabyte 8mm Cartridge Drive with Tape Inserted**

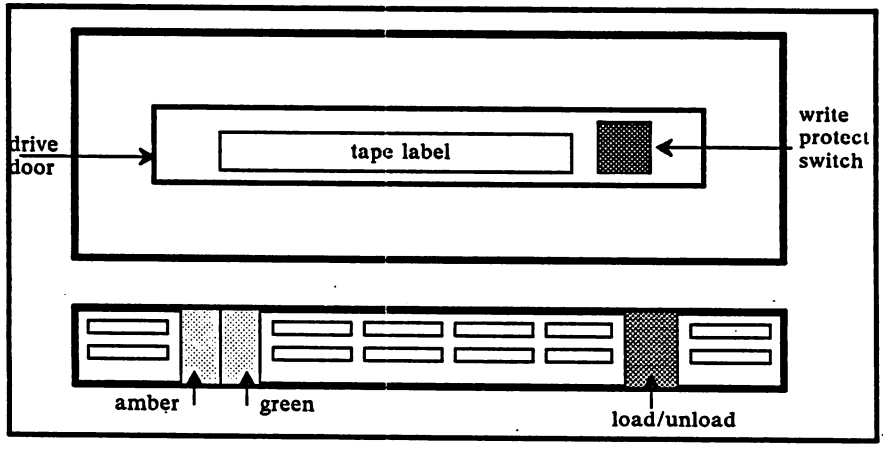


The Alliant 8mm cartridge loads when it receives the first command. When the green LED is on, it indicates the cartridge is loaded and ready.

Variable blinking of the amber LED indicates SCSI bus activity (commands to the drive).

### 3.3.5 Unloading an 8mm Cartridge

To unload an Exabyte 8mm cartridge:

1. Press the LOAD/UNLOAD button on the lower right side of the drive. The green LED turns off. The cartridge tape system rewinds the tape, unloads, and ejects the cartridge.

   If the tape is not automatically unloaded and ejected, check for the following conditions:

   - drive is not powered on
   - drive is not idle
   - contingent connection to/from host
   - mt command in effect (tape cannot be ejected if this command is in effect)

2. Remove the cartridge and close the door.

You can use the mt command as follows to make a loaded device inaccessible:

```
# mt -f /dev/rschs02 offline
```

### 3.3.6 8mm Cartridge Power-up Initialization Indicators

During power-up initialization, the amber and green LEDs both turn on during power-on self-test diagnostics. This lasts about two minutes. When the diagnostics are complete, both LEDs turn off. If the self-test fails, both green and amber LEDs flash.

After initialization, the green LED indicates the status of the drive. When the green light is on, it indicates the drive is loaded and ready.

Variable blinking of the amber LED when the drive is in a ready state indicates the drive is busy.

## 3.4 TAPE DEVICES (9-TRACK)

Alliant currently supports the **Xylogics 772 (VMEbus) tape controller**. This controller supports half-inch tape drives.

### 3.4.1 Tape Names

The conventional names for tape devices (special files) are as follows:

**/dev/rxtv***cud* Xylogics VMEbus

**/dev/nrxtv***cud* Xylogics VMEbus, no rewind

Where:

*c* is an integer 0 – 11 that specifies a tape controller.

*u* is an integer 0 – 7 that specifies a unit (drive).

*d* is a single alphabetic character that specifies density; the options are h (6250bpi), m (1600bpi), l (800bpi), and x (3200bpi).

Each tape drive requires one special file for each density at which it operates. An additional special file is required (at each density) to not rewind the tape after an operation. Typical access to a tape drive is at 6250bpi with rewind, for example, /dev/rxtv00h to access controller 0, unit 0 and /dev/rxtv10h to access controller 1, unit 0.

Density is controlled through software, not switches on the tape drive.

### 3.4.2 Supported 9-Track Tape Drives

The following tape drives, compatible with Xylogics tape controllers, are supported by Concentrix-2800 as shipped:

- **cipher** – At 6250bpi (high speed, short gap), 1600bpi (high speed, short gap), and 3200bpi (high speed, short gap), with automatic back-up after a write error.
- **telex50** – (50ips) At 6250bpi (high speed, short gap), 1600bpi (high speed, short gap), and 800bpi (high speed, short gap).
- **telex125** – (125ips) At 6250bpi (high speed, short gap), 1600bpi (high speed, short gap), and 800bpi (high speed, short gap), with automatic back-up after a write error.

The following tape drives, compatible with Xylogics tape controllers, are defined but not supported:

- **anritsu** – At 1600bpi (high speed, short gap) and 3200bpi (high speed, short gap), with automatic back-up after a write error.
- **fujitsu** – At 6250bpi (high speed, short gap), 1600bpi (high speed, short gap), and 800bpi (high speed, short gap), with automatic back-up after a write error.

- **kennedy9100** – High density (high speed, short gap) and low density (high speed, short gap), with automatic back-up after a write error.
- **kennedy** – At 1600bpi (high speed, short gap) and 800bpi (high speed, short gap), with automatic back-up after a write error.
- **keystone** – At 6250bpi (high speed, short gap) and 1600bpi (high speed, short gap), with automatic back-up after a write error.
- **stc** – At 6250bpi (high speed, short gap) and 1600bpi (high speed, short gap), with automatic back-up after a write error.
- **thorn** – At 6250bpi (high speed, short gap), 1600bpi (high speed, short gap), 800bpi (high speed, short gap), and 3200bpi (high speed, short gap), with automatic back-up after a write error.

## 3.4.3 Configuring Tape Drives

The preferred procedure for configuring a tape drive is to execute the administration command /dev/mktape. This command creates eight special files in /dev with the appropriate major number and minor number to permit access to the given controller and unit (rewind and no rewind at each of four densities).

The following example configures the tape drive on controller 0, unit 0 as a cipher drive, and creates special files to access the drive in 6250bpi, 1600bpi, 800bpi, and 3200bpi:

```
# /usr/alliant/mktape 0 0 cipher
created /dev/rxtv00h
created /dev/nrxtv00h
created /dev/rxtv00m
created /dev/nrxtv00m
created /dev/rxtv00l
created /dev/nrxtv00l
created /dev/rxtv00x
created /dev/nrxtv00x
```

The command removes any existing special files for the specified controller and unit before making the new ones. You should review /dev for any symbolic links to the special files. For example, if /dev/telex points to /dev/rxtv00h and you execute the example shown above, you might want to change the name of the link:

```
# rm /dev/telex
# ln -s /dev/rxtv00h /dev/cipher
```

## 3.4.4 Tape Operations

Use half-inch magnetic tape of standard weight on any size reel up to 2400 feet (10.5 inches). Do not use thin tape (3600 feet).

### 3.4.4.1 Loading a Telex Tape

To load a half-inch tape reel on the standard Telex device shipped by Alliant:

5. Open the door of the tape drive and pull the center of the supply hub so that it snaps out (if it is not already out).

6. Take the tape out of its protective covering. Remove the clamp from the tape reel. Insert a write ring if you want to write to the tape; ensure that a write ring is not inserted if you only want to read from the tape.

7. Seat the tape reel on the supply hub and press it in. Press the center of the hub so that it snaps in.

8. For automatic feed, close the door. For manual feed, drop the end of the tape over the guide to the right of the supply hub, feed the tape slightly, and close the door. Automatic feed usually works for 2400-foot reels. For other sizes, or if the tape refuses to load automatically, use manual feed.

9. For automatic feed, press the LOAD button once. For manual feed, press the LOAD button twice. The load process takes about 30 seconds; wait for the LOAD POINT light to come on. If the LOAD CHECK light comes on instead, the load has failed – try again.

10. Press the ON LINE button. The ON LINE light comes on and the tape unit is ready.

### 3.4.4.2 Unloading a Telex Tape

To unload a tape reel from the standard Telex device shipped by Alliant:

1. Press the RESET button. Press the UNLOAD button. The unload process takes about 30 seconds.

2. Open the door of the tape drive and pull the center of the supply hub so that it snaps out.

3. Pull the tape reel from the supply hub. Close the door.

4. Replace the tape clamp on the reel (without crimping the tape).

### 3.4.4.3 Loading a Cipher Tape

To load a half-inch tape reel on the standard Cipher device shipped by Alliant:

11. Open the door of the tape drive.

12. Take the tape out of its protective covering. Remove the clamp from the tape reel. Insert a write ring if you want to write to the tape; ensure that a write ring is not inserted if you only want to read from the tape.

13. Insert the tape face up (write enable ring down). Ensure that the tape rests firmly on the mounting apparatus.

14. Close the door.

15. Press the LOAD button.

16. Press the ON LINE button. The display panel shows the message *"loading..."* flashing on and off as the tape loads, and the LOAD and ON LINE button lights flash. When the tape is loaded, the lights stop flashing, and stay on.

    The WRT EN (write enable) button's light goes on if the tape is write-enabled. If the tape has been previously written, the tape drive reads the tape's header, and prints the message *"I D <density>"* (density corresponding to the type of tape loaded) on the display panel.

### 3.4.4.4 Unloading a Cipher Tape

To unload a tape reel from the standard Cipher device shipped by Alliant:

1. Press the ONLINE button.

2. Press the UNLOAD button. The message *"unloading..."* is displayed during the unloading process, and the UNLOAD button goes on. When the tape is unloaded, the message *"unloaded"* appears.

3. Open the door of the tape drive and slide the tape out. Close the door.

4. Replace the tape clamp on the reel (without crimping the tape).

### 3.4.4.5 Examining a Tape

Typically, you deal with tapes in afio or tar format for transferring and archiving files, and dump format for dumping and restoring file systems. See Chapter 4. Tapes in other formats require special handling.

Use the /usr/alliant/rt command (read magnetic tape and print record sizes) to discover the number of records on a tape and their sizes. The following command examines /dev/rxtv00h. A record size of 0 indicates an end-of-file marker. Type CTRL C to terminate the operation; otherwise, the operation lasts to the end of the reel, or an error condition.

```
# /usr/alliant/rt /dev/rxtv00h
10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240.
10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240. 10240.
10240. 10240. 0. 0. CTRL c
```

Use the /bin/dd command to read a tape when you do not know the format. First discover the record size, for example, as shown above.

The following command reads from /dev/rxtv00h:

```
# /bin/dd if=/dev/rxtv00h bs=10240
./   775   1765    152          0  4050420012   4215 ./lrm-01   555
1765     40       74511  4031316104   5245

The Ada Joint Program Office does not guarantee the accuracy of this
file, as compared with the contents of ANSI/MIL-STD-1815A-1983,
the Reference Manual for the Ada* Programming Language.  If errors or
discrepancies are found in this machine-readable version, please
forward comments via the Defense Data Network (DDN) to

       Ada-Info@Ada20.ISI.EDU

or via conventional mail to

       Ada Information Clearinghouse
       3D139 (1211 S. Fern, C-107)
       The Pentagon
       Washington, D.C.  20301-3081.
------------------------------------------------------------------------

...
```

Pipe the output of /bin/dd to /bin/od (octal, decimal, hex, ASCII dump) to examine binary tapes:

```
# /bin/dd if=/dev/rxtv10h bs=10240 | /bin/od -c
```

## 3.4.5 Special Tape Drives

The magtape table contains structures for three dummy tape drives named special1, special2, and special3. You can modify these structures to define new types of tape drives for your system.

### 3.4.5.1 Modifying a Structure for a Special Tape Drive

Each tape drive type is defined by a 16-byte structure in /unix that starts at a symbolic address that is the name of the drive type. For example, the structure for special1 starts at the address special1+0x0; the second word of special1 starts at special1+0x2; and so on. Each structure has the following format:

| | | | |
|---|---|---|---|
| byte 0 | density/speed/gap | pass-thru codes | for 6250bpi |
| byte 2 | density/speed/gap | pass-thru codes | for 1600bpi |
| byte 4 | density/speed/gap | pass-thru codes | for 800bpi |
| byte 6 | density/speed/gap | pass-thru codes | for 3200bpi |
| byte 8 | pointer to name of structure | | |
| byte 12 | flags | | |

Specify a density/speed/gap setting and pass-thru codes for each density at which the tape drive operates. Specify the two bytes as 0 for densities that are not used, except that the 6250bpi density must be specified – if the tape drive does not support 6250bpi, duplicate one of the other specifications, typically the one for the highest density. The density/speed/gap setting is a sum of the following values:

| | | |
|---|---|---|
| 0x2 | low density | choose 1; use only in alternate mode |
| 0x3 | high density | |
| 0x4 | 800bpi | choose 1; use in standard mode |
| 0x5 | 1600bpi | |
| 0x6 | 3200bpi | |
| 0x7 | 6250bpi | |
| 0x10 | high speed | |
| 0x20 | short gap | |

Sum the pass-thru codes that you use:

| | |
|---|---|
| 0x1 | reverse line |
| 0x2 | write line |
| 0x4 | edit line |
| 0x8 | write-file-mark line |
| 0x10 | erase line |
| 0x80 | long gap line |

Do not change the pointer to the name of the structure (bytes 8–11).

Specify the flags as the sum of the following:

| | |
|---|---|
| 0x1 | automatic back-up of one record if write error occurs |
| 0x2 | alternate mode (use low/high density) |

The following example (with italicized annotations) examines the structure for a standard telex50 tape drive:

```
# adb /unix -
ready
telex50?x

_telex50:       0x1d37          short gap, high speed, 6250bpi
                                erase, write-file-mark, edit, reverse
_telex50+0x2:   0x1c35          short gap, high speed, 1600bpi
                                erase, write-file-mark, edit
_telex50+0x4:   0xffff9c34      short gap, high speed, 800bpi
                                long gap, erase, write-file-mark, edit
_telex50+0x6:   0x0             does not support 3200bpi

_telex50+0x8:   0xffffc098      pointer to "telex50"

_telex50+0xA:   0xffffcce

_telex50+0xC:   0x0             no automatic back-up

_telex50+0xE:   0x0
CTRL D
```

To define a new tape drive, write to special1, special2, or special3 in /unix with adb. (Changing the standard structures is not recommended.) The following example defines special1. Be sure to use lowercase w to write to a word and uppercase W to write to a longword. Be sure not to write over the pointer to the structure name at offset 0x8.

```
# adb -w /unix -
ready
special1?w 0x31d37
_special1:      0x0     =       0x1d37
special1+0x2?w 0x1c35
_special+0x2:   0x0     =       0x1c35
special1+0x4?w 0x9c34
_special1+0x4:  0x0     =       0x9c34
special1+0x6?w 0x1c36
_special1+0x6:  0x0     =       0x1c36
special1+0xc?W 0x1
_special1+0xC:  0x1             =       0x1
CTRL D
```

### 3.4.5.2 Configuring Special Tape Drives

You can configure a special tape drive with /usr/alliant/mktape. The following example configures the tape drive on controller 1, unit 0 as a special1 drive, and creates special files to access the drive in 6250bpi, 1600bpi, 800bpi, and 3200bpi:

```
# /dev/mktape 1 0 special1
created /dev/rxtv10h
created /dev/nrxtv10h
created /dev/rxtv10m
created /dev/nrxtv10m
created /dev/rxtv10l
created /dev/nrxtv10l
created /dev/rxtv10x
created /dev/nrxtv10x
```

To make the special files with /etc/mknod, use the following form:

```
# /etc/mknod /dev/[n]rxtvcud c 0x80 0xccuubd
```

The variable portion of the device name specifies the controller, unit, and density, as explained at the beginning of this section. The n specifies no rewind.

The access mode must be c for character. The major number must be 0x80 (or 128 decimal). The 24-bit minor number is specified as follows:

**Rules for Specifying 24-bit Minor Number**

| name | bits | specifies |
|---|---|---|
| *cc* | 23:16 | Controller number |
| *uu* | 15:8 | Unit number |
| *b* | 7:4 | Brand of tape drive:<br>0x0 – telex50 0x5 – kennedy9100 0xa – thorn<br>0x1 – cipher 0x6 – pertec 0xb – special1<br>0x2 – fujitsu 0x7 – kennedy 0xc – special2<br>0x3 – telex125 0x8 – anritsu 0xd – special3<br>0x4 – keystone 0x9 – stc |
| *d* | 3:0 | Density (3:1) and Rewind (0):<br>0x0 – 6250 bpi rewind 0x1 – 6250,bpi no rewind<br>0x2 –1600 bpi, rewind 0x3 – 1600 bpi, no rewind<br>0x4 – 800 bpi, rewind 0x5 – 800 bpi, no rewind<br>0x6 – 3200 bpi, rewind 0x7 – 3200 bpi, no rewind |

## 3.5 TERMINAL DEVICES

Alliant currently supports the **Xylogics 780 (VMEbus)** terminal multiplexor. This terminal multiplexor supports up to 16 RS232 devices.

The VMEDP contains a duart (dual universal asynchronous receiver transmitter) for terminal and modem connections.

### 3.5.1 Terminal Names

The conventional names for devices connected to the VMEDP duarts are as follows. In all cases, the major number is 0xc1 (193).

| | |
|---|---|
| **/dev/console** | device connected to the console port on IP0; minor number 0 |
| **/dev/ttyR0** | device connected to the remote port on IP0; minor number 1 |
| **/dev/ttyC*u*** | additional console ports |
| **/dev/ttyR*u*** | additional remote ports |

The conventional name for a device connected to a multiplexor port is as follows. In all cases, the major number is 0x98 (152):

**/dev/tty*cu***

Where:

| | |
|---|---|
| *c* | identifies the multiplexor. |
| *u* | identifies the multiplexor ports (specify *u* as a hexadecimal digit). |

The minor number for a multiplexor device is specified as follows in hexadecimal:

| | |
|---|---|
| 0x*cc*000*u* | for a terminal (bit 7 is clear) |
| 0x*cc*008*u* | for a modem (bit 7 is set) |

For example, /dev/tty0e has one of the following minor numbers:

| | |
|---|---|
| 0x00000e | terminal |
| 0x00008e | modem |

To enable access to a unit as either a terminal or modem, create two special files: one with bit 7 of the minor number cleared and one with bit 7 set. Use the conventional name for the more common occurrence and a suffixed name for the alternate occurrence. For example:

| | terminal (bit 7 = 0 ) | modem (bit 7 = 1) |
|---|---|---|
| usually terminal | /dev/tty*cu* | /dev/tty*cu*.modem |
| usually modem | /dev/tty*cu*.nomodem | /dev/tty*cu* |

In addition, the device /dev/tty (major number 0xc2 (194), minor number 0) refers to the terminal of the current process.

### 3.5.1.1 Script mknewdevs

The script /dev/mknewdevs, by default, makes special files for the following devices (console, process tty, remote port, and 16 terminal devices attached to multiplexor 0):

```
crw-rw-rw-  1 root     wheel    0xc1,      0 May 21 15:32 console
crw-rw-rw-  1 root     wheel    0xc2,      0 May 21 15:32 tty
crw-rw-rw-  1 root     wheel    0xc1,    0x1 May 21 15:32 ttyR0
crw-rw-rw-  1 root     wheel    0x98,      0 May 21 15:33 tty0
crw-rw-rw-  1 root     wheel    0x98,    0x1 May 21 15:33 tty1
crw-rw-rw-  1 root     wheel    0x98,    0x2 May 21 15:33 tty2
crw-rw-rw-  1 root     wheel    0x98,    0x3 May 21 15:33 tty3
crw-rw-rw-  1 root     wheel    0x98,    0x4 May 21 15:33 tty4
crw-rw-rw-  1 root     wheel    0x98,    0x5 May 21 15:33 tty5
crw-rw-rw-  1 root     wheel    0x98,    0x6 May 21 15:33 tty6
crw-rw-rw-  1 root     wheel    0x98,    0x7 May 21 15:33 tty7
crw-rw-rw-  1 root     wheel    0x98,    0x8 May 21 15:33 tty8
crw-rw-rw-  1 root     wheel    0x98,    0x9 May 21 15:33 tty9
crw-rw-rw-  1 root     wheel    0x98,    0xa May 21 15:33 ttya
crw-rw-rw-  1 root     wheel    0x98,    0xb May 21 15:33 ttyb
crw-rw-rw-  1 root     wheel    0x98,    0xc May 21 15:33 ttyc
crw-rw-rw-  1 root     wheel    0x98,    0xd May 21 15:33 ttyd
crw-rw-rw-  1 root     wheel    0x98,    0xe May 21 15:33 ttye
crw-rw-rw-  1 root     wheel    0x98,    0xf May 21 15:33 ttyf
access mode                     major    minor             name
```

To change the initialization default to modem devices, edit the script /dev/mknewdevs. Edit the following line, substituting an "m" for the "t" at the end of the line:

```
/dev/mkttynod xmv 0 t          terminals;default
/dev/mkttynod xmv 0 m          modems
```

Add mkttynod lines to create additional multiplexor devices with the mknewdevs script.

### 3.5.1.2 Command mkttynod

The administration command /dev/mkttynod makes special devices for a set (one controller) of 16 ports. The following example demonstrates the addition of special files for controller 1. For the first argument, specify xmv (VMEbus); the second argument is the controller number; specify m instead of t for modem devices.

```
# /dev/mkttynod xmv 1 t
```

### 3.5.1.3 Command mknod

To add special files on an individual basis or to use nonbase names, use the administration command /etc/mknod. The form of the command is as follows:

```
# /etc/mknod /dev/ttycu c 0x98 0xcc00mu
```

The variable portion of the device name specifies the controller and port number. The access mode must be c for character. The major number for a terminal is 0x98 (152). The 24-bit minor number is specified as follows:

**Rules for Specifying the Minor Number for a Terminal**

| name | bits | specifies |
|---|---|---|
| *cc* | 23:16 | Controller number |
| | 15:8 | Not used |
| *m* | 7 | Modem control bit:<br>0 = no modem<br>1 = modem |
| | 6:4 | Not used |
| *u* | 3:0 | Port number |

The following command makes a special file for a modem device:

```
# /etc/mknod /dev/tty1f.modem c 0x98 0x01008F
```

#### 3.5.1.4 Pseudo Terminals

See Chapter 7 for a discussion of pseudo terminals.

## 3.5.2 Hardware Configuration

Typically, the back of the system cabinet contains two duart ports for the console and the modem (remote port); and the back of the system or I/O expansion cabinet contains 16, or 32 terminal ports (16 per controller). The following figure demonstrates a typical arrangement. The groups of eight may be positioned differently but the numbering within each group is as shown.

**Terminal Ports**

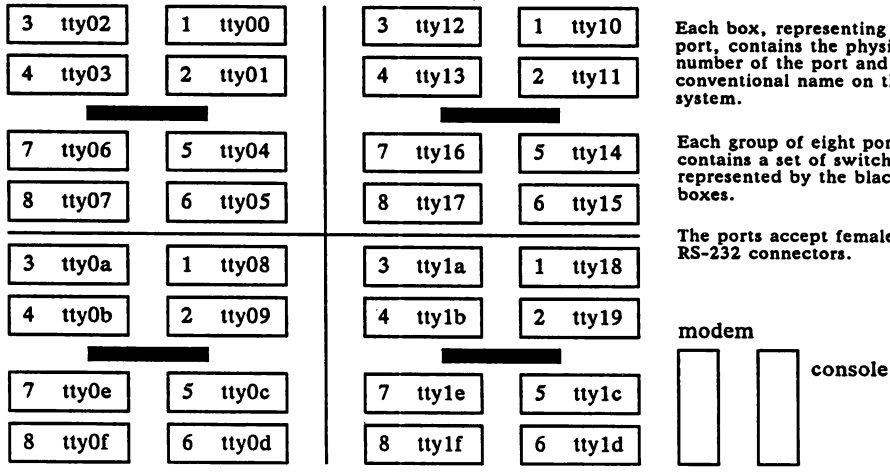

### 3.5.2.1 Cabling

Each 8-port group on the **VMEbus multiplexor** contains two 60-pin headers for receiving the ribbon cable from the controller proper. The multiplexor (for those ports) operates as a DCE device if the ribbon cable is connected to the upper header and as a DTE device if the ribbon cable is connected to the lower header. DCE mode permits use of the standard Alliant cables; DTE mode permits the use of straight-through cables for modem connections. The standard mode is DCE; you must remove the back panel of the machine to access the headers, which are on the inside of the port panel.

The following tables describe the pin connections used by Alliant equipment.

**Alliant pin connections**

| pin | function | DCE | DTE |
|---|---|---|---|
| 1 | chassis ground | - | - |
| 2 | TXD | in | out |
| 3 | RXD | out | in |
| 4 | RTS | in | out |
| 5 | CTS | out | in |
| 6 | DSR | out* | in* |
| 7 | signal ground | - | - |
| 8 | DCD | in* | in* |
| 20 | DTR | in* | out* |

| | |
|---|---|
| TXD | transmit data |
| RXD | receive data |
| RTS | request to send |
| CTS | clear to send |
| DSR | data set ready |
| DTR | data terminal ready |
| DCD | data carrier detect |

* On a VMEbus multiplexor, exactly one of these three lines can be monitored as an input. The other two are ignored. Monitoring an output signal (out*) is possible, but not useful.

The following diagrams describe the standard Alliant cables. These connections treat Alliant ports as DCE devices and monitor DCD on pin 8. These cables require that a VMEbus multiplexor be in DCE mode with the DCD switch closed.

- **Standard terminal connection** – This cable connects a terminal or data switch to an Alliant multiplexor port, and connects a terminal to the console port.

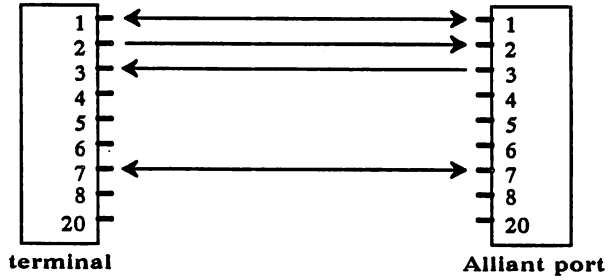


- **Standard modem connection** – This cable connects a modem to an Alliant multiplexor port or to the remote port.

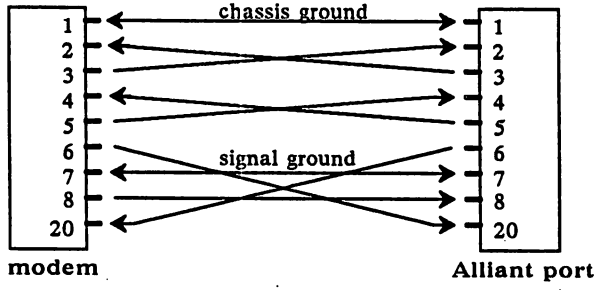


- **Standard console printer connection** – This cable connects the console to the dot matrix printer.

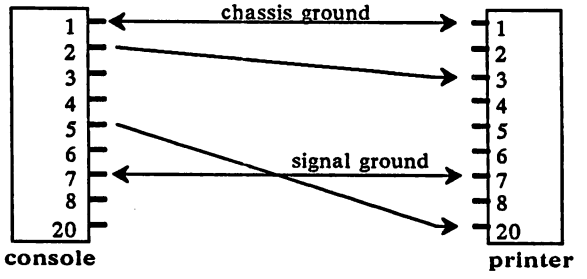

### 3.5.2.2 Switches

The VMEbus multiplexor has four sets of eight switches for each group of ports, running right to left and numbered 1– 8 left to right. In set 1, each switch corresponds to the port with the same number. In the remaining three sets (24 switches), each group of three adjacent switches corresponds to the next port starting at port 1.

**VMEbus multiplexor switches**

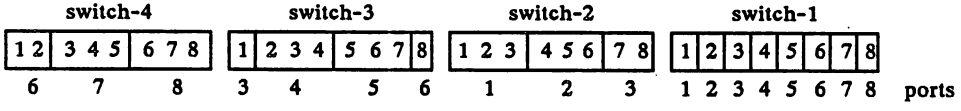


Setting a switch-1 on (closing the switch) shorts CTS (pin 5) to RTS (pin 4) on the corresponding port. For Alliant standard cables, the switch should be off to connect to a modem and on to connect directly to a terminal.

For switch-2, switch-3, and switch-4, set one switch in each group on (closed) as follows (one switch must be on, two must be off):

- **DCD (leftmost)** – To monitor DCD (pin 8).
- **DSR (middle)** – To monitor DSR (pin 6).
- **DTR (rightmost)** – To monitor DTR (pin 20).

The typical setup for an Alliant VME multiplexor is always the same. The following switch settings permit hardware modem control to be used on any multiplexor port. (Modem control must still be enabled in the software by setting the port's minor device number appropriately).

**VMEbus multiplexor switches**

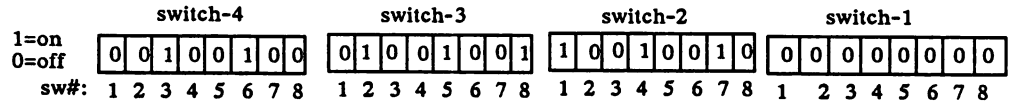


This setup monitors DCD on all ports, and does not short RTS to CTS. If modem control is enabled on a port, Concentrix will send a "hangup" (SIGHUP) signal to the appropriate processes if the monitored signal (DCD in the example) is lost (if the user hangs up). Concentrix will drop DSR when a user logs out. Thus, for a properly configured modem, two things should be true:

- If a user logs out, the modem should hang up. (The same will happen if the user does not enter a correct ID/password pair quickly enough.)
- If a user hangs up without logging out, Concentrix should log the user out automatically. This is a particularly important security feature for dialup lines.

If this handshaking fails, check for the following:

- Incorrect wiring
- Incorrect setting of VME multiplexor switches
- Minor device number on TTY port (/dev/tty??) < 128
- The modem itself is "faking" DCD, or is not recognizing the request to hang up. With the standard wiring configuration, the modem should provide DCD (data carrier detect) only when an actual carrier is present; it should also hang up if it detects that DTR (data terminal ready) has been dropped. In many current modems, these are configurable options which may not have been enabled at the factory.

## 3.5.3 Software Configuration

The file /etc/ttys (see the *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ttys) controls terminal connections. Each line of the file specifies one terminal as follows:

- **Terminal name (column 1)** – The terminal name starts in column 1. Use only the file name, not the full pathname.
- **Control program (column 2)** – Column 2 specifies a command line that invokes a program to initially control the terminal. The typical command is /etc/getty with one argument; the argument identifies a set of terminal characteristics described in the file /etc/gettytab (see below). The command line must be enclosed in quotation marks if it contains white space.
- **Terminal type (column 3)** – Column 3 specifies the type of terminal usually connected as the device. The terminal type must be a valid type specified in the file /etc/termcap (see the *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: termcap).

The remaining columns specify flags and (optionally) a window control program:

- **On flag** – The keyword on activates the terminal for logins, that is, invokes the control program for the program. No on flag or the keyword off deactivates the terminal, that is, does not invoke the control program. Specify on for direct terminals and in-coming modems unless you have a special reason to deactivate the terminal. Specify off for outgoing modems and printers.
- **Secure flag** – The keyword secure permits the super-user (root) to log in on the terminal. Omit this flag to prevent super-user login from a terminal.
- **Window control program** – The keyword window= introduces a command line that invokes a windowing program that is executed prior to the control program. The command line must be enclosed in quotation marks.

The typical sets of terminal characteristics, taken from the file /etc/gettytab (see the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: gettytab) as distributed, are as follows:

- **0** – (aliases: d300, Dial-300) Dial-in line that cycles between 300, 1200, 150, and 110 baud when the user presses BREAK.
- **2** – (aliases: std.9600, 9600-baud) Direct-connect 9600-baud line.
- **3** – (aliases: D1200, Fast-Dial-1200) Dial-in line that cycles between 1200, 300, and 2400 baud when the user presses BREAK.
- **5** – (aliases: D300, Fast-Dial-300) Dial-in line that cycles between 300, 2400, and 1200 baud when the user presses BREAK.
- **h** – (aliases: h2400, High-Speed-Dialup-2400) Dial-in line that cycles between 2400, 1200, and 300 baud when the user presses BREAK.
- **m** – (aliases: m9600, High-Speed-Direct-9600) Direct-connect line that cycles between 9600, 2400, 1200, and 300 baud when the user presses BREAK.

The file /etc/ttys as distributed contains an entry for the console. Entries for tty00 through tty0f are present but commented out. Entries for ttyp0 through ttypf and ttyq0 through ttyqf are present but specify the control program as none.

You must uncomment and/or adjust the existing entries, and add entries, to suit your configuration. The following example demonstrates /etc/ttys entries for the console and 16 terminals connected to controller 0. Half the lines are direct at 9600 baud and half the lines are dial-in at 1200-300 baud. The direct lines permit superuser login, the dial-in lines do not.

```
console "/etc/getty std.9600" vt100   on secure
tty00   "/etc/getty std.9600" vt100   on secure
tty01   "/etc/getty std.9600" vt100   on secure
tty02   "/etc/getty std.9600" vt100   on secure
tty03   "/etc/getty std.9600" vt100   on secure
tty04   "/etc/getty std.9600" vt100   on secure
tty05   "/etc/getty std.9600" vt100   on secure
tty06   "/etc/getty std.9600" vt100   on secure
tty07   "/etc/getty std.9600" vt100   on secure
tty08   "/etc/getty 3"        vt100   on
tty09   "/etc/getty 3"        vt100   on
tty0a   "/etc/getty 3"        vt100   on
tty0b   "/etc/getty 3"        vt100   on
tty0c   "/etc/getty 3"        vt100   on
tty0d   "/etc/getty 3"        vt100   on
tty0e   "/etc/getty 3"        vt100   on
tty0f   "/etc/getty 3"        vt100   on
```

The standard Alliant terminals have a terminal type of vt100. Note that the termcap file supplied with Concentrix-2800 includes terminal descriptions for a number of AT&T terminals.

To disable a terminal while the system is running, edit the file /etc/ttys and remove the on flag of the appropriate line (or change it to off). Then send a hangup signal to the init process by typing:

```
# kill -HUP 1
```

To enable a configured terminal while the system is running, edit the file /etc/ttys and add the on flag to the appropriate line (or change off to on). Then send the hangup signal to the init process as shown above.

If a terminal is inaccessible when init tries to create a process for it, init prints a message on the console and tries to reopen the terminal every minute, reprinting the warning message every 10 minutes.

### 3.5.3.1 Logging In from An Apollo Window

Alliant provides a script to properly configure an Apollo window as a terminal on a remote login to Concentrix-2800. If the remote login is your usual access mode to Concentrix-2800, you should place the following line in your .login file:

```
source /usr/alliant/apollo.vt100.setsize
```

See also Add_terminal in Appendix B.

## 3.6 SYSTEM MEMORY MAPS

The following tables demonstrate the use of memory by Concentrix-2800 on CEs and IOMs, and by Diagnostix-2800 on the Diagnostic IP:

The Concentrix-2800 CE/IOM virtual address space layout is subject to change.

Sizes of segments and pages are:

- One segment = 4MB (0x400000)
- One page = 4KB (0x1000)

The following table shows the Concentrix-2800 virtual memory address ranges for various system sections, and the number of segments reserved for each. Where applicable, the kernel variables' symbols are supplied.

**Concentrix-2800 Memory Map**

| address range | description | # segs | symbol | Area* |
|---|---|---|---|---|
| ffffffff-ffc00000 | Processor private area | 1 | gpt_urodatava | K |
| ffbfffff-ff800000 | Process private area | 1 | ppt_privateva | K |
| ff7fffff-ff400000 | Process page tables | 1 | Sysmap | K |
| ff3fffff-ff000000 | Ddt image area | 1 | ddtva | K |
| feffffff-fd000000 | HW registers, VMEbus | 8 | gpt_hardva | K |
| fcff6fff-fcc00000 | Unix image | 1 | btext | K |
| fcbfffff-fc800000 | RT driver | 1 | gpt_rtdriverva | K |
| fc7fffff-e8800000 | global data area | 80 | | K |
| e87fffff-d8800000 | mapped file area | 64 | | K |
| d87fffff-d6800000 | imuac area | 8 | | K |
| d67fffff-variable | user arguments | 128KB | | U |
| variable-variable | user stack | 6 | | U |
| variable-variable | user cactus stack | 8 | | U |
| variable-00800000 | user code and data | 3.25GB | | U |
| 00799999-00001000 | shared libraries | 2 | | U |
| 00000fff-00000000 | nonexistent memory | 4KB | | |

* K = Kernel area; U = User area

The following tables give further refinement, in pages, of four of the system sections listed above.

**Concentrix-2800 Memory Map: Processor Private Area**

| address | description | # pages | symbol |
|---|---|---|---|
| ffff6000 | Cache flush pages | 2 | ppt_flushva |
| ffc18000 | unused | 10 | |
| ffc17000 | ddt debugger | 1 | ppt_ddtdebugva |
| ffc16000 | Alliant | 1 | |
| ffc06000 | Alliant | 16 | |
| ffc03000 | system read/write data | 3 | ppt_rwdatava |
| ffc02000 | nonexistent memory | 1 | |
| ffc01000 | system readonly data | 1 | ppt_rodatava) |
| ffc00000 | user readonly data | 1 | gpt_urodatava |

**Concentrix-2800 Memory Map: Process Private Area**

| address | description |
|---|---|
| ff879000 | other dynamically mapped locations, subject to change |
| ff805000 | kernel stacks, 29 of them @ 3 pages each (1 for each processor) |
| ff804000 | nxm |
| ff800000 | "u" area, currently 4 pages |

**Concentrix-2800 Memory Map: Hardware Area**

**(16 regions of 2MB each (8 segs))**

| address range | description |
|---|---|
| feffffff-fee00000 | IMB ring buffers (low in VME 0 or 1 standard space) |
| fedfffff-fec00000 | dyn. allocated to VME 0 or 1 standard or extended space |
| febfffff fea00000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fe9fffff-fe800000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fe7fffff-fe600000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fe5fffff-fe400000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fe3fffff-fe200000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fe1fffff-fe000000 | dyn. allocated to VME 0 or 1 standard or extended space |
| fdffffff-fde00000 | For VMEbus 1 interrupt acknowlege cycles |
| fddfffff-fdc00000 | For VMEbus 0 interrupt acknowlege cycles |
| fdbfffff-fda00000 | VMEAC 1 (or other adapter card) registers |
| fd9fffff-fd800000 | VMEAC 0 (or other adapter card) registers |
| fd7fffff-fd600000 | IOM i/o maps (256K repeated through 2MB) |
| fd5fffff-fd400000 | IOM registers |
| fd3fffff-fd200000 | VMEbus 1 short space (64K repeated through 2MB) |
| fd1fffff-fd000000 | VMEbus 0 short space (64K repeated through 2MB) |

**Concentrix-2800 Memory Map: Global Area**

| address range | description | # segs | symbol |
|---|---|---|---|
| fc7fffff-fc400000 | Kernel static allocate | 1 | valloc |
| fc3fffff-fbc00000 | Kernel dyn. allocation | 2 | kmem |
| fbbfffff-ebc00000 | Dyn. user page table | 64 | upgtbls |
| ebbfffff-e9c00000 | Core map | 8 | cmap |
| e8cfffff-e8c00000 | Buffer cache | 4 | buf |
| e8bfffff-e8800000 | User interrupt service | 1 | isr |

The following table provides further refinement of the user readonly data page in the Processor Private area.

**Concentrix-2800 Memory Map: Processor Private Area - gpt_urodatava page**

| address | description | symbol |
|---|---|---|
| ffc00028 | number of cpus in system | sys_number_of_cpus |
| ffc00024 | system model number | modelnum |
| ffc00020 | system serial number | serialnum |
| ffc00018 | chip instruction cache size | onchip_icachesize |
| ffc00014 | chip data cache size | onchip_dcachesize |
| ffc00010 | global cache size | global_cachesize |
| ffc0000c | system line clock hertz rate | sys_hz |
| ffc00008 | user readonly idle | uro_idle_addr |
| ffc00000 | hardware counter | hrcounter |

The following table shows the virtual and physical memory map for Diagnostic-2800.

**Diagnostix-2800 Memory Map**

| virtual | description | | size | physical |
|---|---|---|---|---|
| 0xffffffff | | | | |
| 0xffc00000 | unused (so small negative pointers will buserr) | | 4MB | |
| 0xffb00000 | unused | | 1MB | |
| 0xffa00000 | unused | | 1MB | |
| 0xff900000 | unused | | 1MB | |
| 0xff800000 | unused | | 1MB | |
| 0xff000000 | unused | | 8MB | |
| 0xfe000000 | VME standard addresses | | 16MB | 0x98000000 |
| 0xf0000000 | unused | | 224MB | |
| 0xe0000000 | VME extended asddresses | | 256MB | 0x80000000 |
| 0x00800000 | unused (general) | | 3576MB | |
| 0x00782000 | unused (for page table expansion) | | 504KB | |
| 0x00780000 | PROM generated page tables | 2048L = | 8KB | 0xa03fb000 *(near end of local DRAM)* |
| 0x00771000 | unused (for segment table expansion) | | 60KB | |
| 0x00770000 | PROM generated segment tables | 1024L = | 4KB | 0xa03fd000 *(near end of local DRAM)* |
| 0x0076f000 | function code map 8L + 1016L = | | 4KB | 0xa03fe000 *(near end of local DRAM)* |
| 0x00704000 | unused(for I/O map expansion) | | 428KB | |
| 0x00700000 | VME-to-RAM I/O maps | 4096*4 = | 16KB | 0x98000000 *(VME standard addresses)* |
| 0x006f0000 | VME short addresses | | 64KB | 0x90000000 |
| 0x006e8000 | EPROM (device) | | 32KB | 0xc0008000 |
| 0x006e0000 | EPROM (power) | | 32KB | 0xc0000000 |
| 0x006df000 | DUART | 128*32 = | 4KB | 0xcc000000 |
| 0x006de000 | system timer | 1024*4 = | 4KB | 0xc4000000 |
| 0x006dd000 | IP contol register | 128*32 = | 4KB | 0xd0000000 |
| 0x006dc000 | boot parameter page | | 4KB | 0xa03ff000 *(local DRAM, last page)* |
| 0x006db000 | SCSI interface | 2048*2 = | 4KB | 0xc8000000 |
| 0x006da000 | DES interface | 2048*2 = | 4KB | 0xe0000000 |
| 0x006d9000 | IP status register | 1024*4 = | 4KB | 0xd4000000 |
| 0x006d8000 | diagnostic register | 1024*4 = | 4KB | 0xdc000000 |
| 0x006d7000 | IPC cache | 1024*4 = | 4KB | 0xd8000000 |
| 0x006d1000 | spare, invalid | | 24KB | |
| 0x006d0000 | IPC device registers | | 4KB | 0x40000000 |
| 0x006c0000 | test IPC dirty bits | | 64KB | 0x400c0000 |
| 0x006b0000 | clear IPC dirty bits | | 64KB | 0x400e0000 |
| 0x006a0000 | set IPC dirty bits | | 64KB | 0x400f0000 |
| 0x00690000 | IPC cross processor interrupt | | 64KB | 0x40080000 |
| 0x00670000 | reserved (DDT is mapped here) | | 128KB | |
| 0x0066ffff | unused | | 2516KB | |
| 0x003fb000 | end local DRAM | | | |
| 0x00000000 | start local DRAM | | 4076KB | 0xa0000000 |

# CHAPTER 4
# FILE SYSTEMS

### Administration Files

| name | description |
|---|---|
| /etc/fstab | static information about file systems ** (edit) |
| /etc/group | group file ** (edit) |
| /etc/mtab | mounted file system table ** |
| /etc/passwd | password file ** (edit) |
| quotas | quota file at the file system root |

** See *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

### Administration Commands (Appendix A)

| name | description |
|---|---|
| automount | automatically mount NFS fiule system |
| dump | incremental file system backup |
| dumpfs | dump file system information |
| edquota | edit user quotas |
| fsck | file system consistency check / interactive repair |
| makefs | construct a file system |
| mklost+found | make lost+found directory for fsck |
| mknod | build a special file |
| mkstripe | make striped partition |
| mount | mount file systems |
| quot | summarize file system usage |
| quotacheck | file system quota consistency checker |
| quotaon | turn file system quotas on and off |
| repquota | summarize quotas for a file system |
| restore | incremental file system restore |
| sync | update the superblock |
| tunefs | tune up an existing file system |
| umount | dismount file systems |
| unmkstripe | disassemble striped partition |
| update | periodically update the superblock |

**Administration Procedures (Appendix B)**

| name | description |
|---|---|
| Format_disk | format and label disk |
| Label_disk | label (relabel) disk |
| Make_file_system | make and mount a new file system |

**General Commands (Concentrix-2800 Commands and Applications Manual)**

| name | description |
|---|---|
| afio | manipulate and archive file system |
| bfio | manipulate and archive file system |
| chgrp | change file group |
| chmod | change file access mode |
| df | report disk free space |
| ls | list contents of directory |
| mkdir | create directory |
| quota | display disk usage and limits |
| tar | tape archiver |

File systems are hierarchical arrangements of directories and ordinary files. Directories can contain other directories and ordinary files; ordinary files are end nodes in the hierarchy. A file system resides in one or more disk partitions and is associated with a parent directory through a process called mounting. The following file systems exist. (Note that the swapping/paging area is not a file system.)

- **root** – The root file system contains the system kernel (the base operating system) and other system software. The root file system is mounted on the root directory, the parent of all files. The disk partition containing the root file must be on a controller connected to IP0.

- **usr** – The usr file system contains additional system software. The usr file system is mounted on the usr directory in the root file system and forms a logical extension to the root file system.

- **other** – You can make and mount other file systems on directories in the root file system. The directory at which a file system is mounted is called the file system mount point or file system root. The mounted file system forms a logical extension to the existing hierarchy starting at the mount point.

The root file system and all mounted file systems are accessed as a single system of files with one exception. No hard link (ln command) can exist between two file systems. Symbolic links (ln -s) can exist between file systems.

The pathname for the root directory is the slash (/). The following example lists the contents of the root directory in long format:

```
# ls -l /
```

The absolute pathname for a directory or file immediately under the root directory is the slash (/) followed by the name of the directory or file. A directory immediately under the root directory may or may not be the root of a file system depending on whether a file system is mounted there. The following example lists the contents of the /usr directory in long format:

```
# ls -l /usr
```

The user command df lists the file systems mounted on your system and provides the following information about each: block device name, size in kilobytes, kilobytes in use, kilobytes available, the percentage of space in use, and the mount point. The following example lists all file systems mounted on your system:

```
# df
```

The following example lists the root and /usr file systems:

```
# df / /usr
```

You can list the names, mount points, and types of file systems currently mounted by typing the mount command without arguments:

```
# mount
```

## 4.1 TYPES OF FILE SYSTEMS

Concentrix-2800 supports the following types of file systems:

- **4.3BSD** – The basic unit of allocation for files is 4096 bytes. User programs access data through system buffers. The file system resides in one disk partition.

- **Track-allocated** – The basic unit of allocation for files is a full track. The actual size of a track varies among types.The track size is 40,960 bytes for fj690s disks, 49,728 bytes for fj1000s disks, and 49,152 bytes formatted for fj2600s disks. The larger unit size results in fewer reads and writes per I/O request and less disk fragmentation. User programs access data directly through the virtual memory system, and asynchronous read-ahead and write-behind provide additional speed. The file system resides in one disk partition.

- **Striped** – A single track-allocated file system is distributed over multiple disks so that, for example, I/O requests on a single file can be performed by several disk controllers on several busses working concurrently. The file system resides in two or more disk partitions.

Track-allocated and striped file systems maximize system I/O throughput but require more disk space (typically 20 per cent more) than 4.3BSD file systems. They work best as file systems with a few large files or with very high usage.

> **Note:** The root and usr file systems must be 4.3BSD file systems.

## 4.2 MAKING NEW FILE SYSTEMS

The root disk includes two root file systems: One for Concentrix-2800 and one for Diagnostix-2800, as explained in Chapter 3. The discussion that follows deals with Concentrix-2800 file systems only.

Typically, a system uses the root disk for the root file system (/), the usr file system (/usr), and the swapping/paging area, and the remaining disks are free for user file systems. This configuration may be modified to enlarge the swapping/paging area; see Section 3.2.4.3. The user file systems can be configured in various ways:

- **Default partitions** – You can simply use each h partition of each remaining disk for one user file system.

- **Default partitions (variation)** – You can put three file systems on one disk by using the r, s, and f partitions.

- **User partitions** – You can relabel the disk to tailor partitions to the number and sizes of file systems needed. See Label_disk in Appendix B, "Adding User Partitions."

- **Disk stripes** – You can put one file system on multiple disks by creating a disk stripe. For concurrent I/O, the partitions should be on multiple disks and controllers.

By default, the directory /tmp is in the root file system. This directory tends to hold large amounts of data and can cause the root file system to fill. The recommended procedure is to create a track-allocated file system in at least an s size partition on a disk other than the root disk and to mount that file system at /tmp. Alternatively, you can make a directory in a file system other than the root file system and make a soft link from /tmp to the new directory. A separate file system is preferable because the /tmp directory in root acts as a fallback in case the the /tmp file system is not mounted; and because the /tmp system might be implemented as a disk stripe to improve throughput. You should examine /tmp periodically and delete unneeded files.

The following figure illustrates some file system configurations. The figure assumes a system with four disks mounted on controllers 0 and 1, units 0 and 1, with standard VMEbus names. Partitions q and p are Diagnostix-2800 root (p) and paging/swapping (q) partitions. They are reserved by Alliant; they can only be used from Diagnostix-2800.

**File System Configurations (Concentrix-2800)**

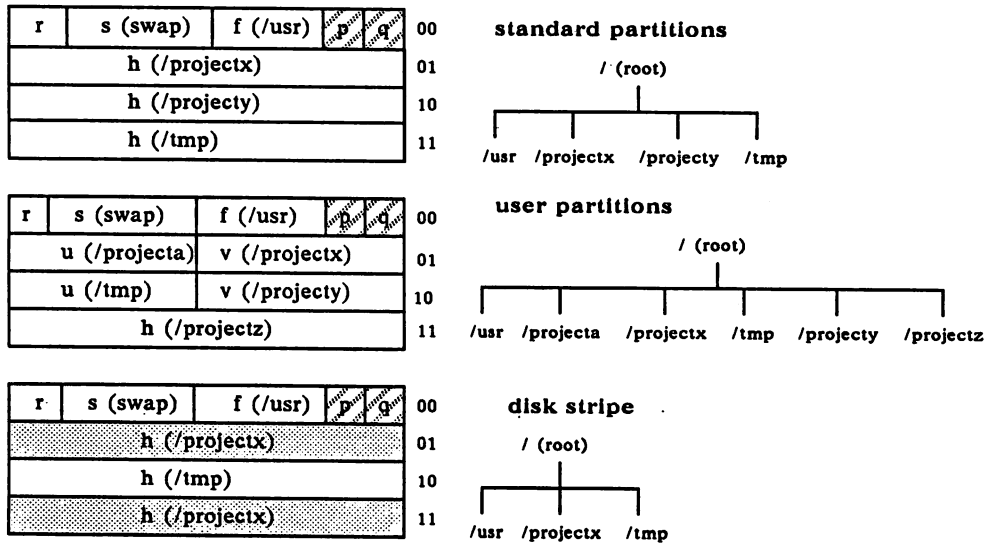


The basic steps in making a new file system are as follows:

1. Make a file system (/etc/makefs command) in the desired partition. Be sure that no file system already exists in that partition.
2. Make a directory (mkdir command) in the root file system to serve as the mount point for the new file system.
3. Mount the file system (/etc/mount command) at the mount point and update the file system table /etc/fstab.

Use of non-standard partitions requires that you first relabel the disk. Use of a disk stripe requires making special files for the striped device and associating the striped device with the participating partitions.

See Make_file_system for detailed procedures. See the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: fstab, for the format of /etc/fstab.

The following example makes a new file system in partition h of disk 01. The new file system starts at /projectx in the root directory.

```
# /etc/makefs -t tfs /dev/rxdv01h
# /bin/mkdir /projectx
# /etc/mount /dev/xdv01h /projectx
# /bin/cp /etc/fstab /etc/fstab.old
# /etc/mount -p > /etc/fstab
```

The next example creates a striped file system for /projectx in the h partitions of disks 01 and 11.

```
# /etc/mknod stripe0 b 200 0
# /etc/mknod rstripe0 c 200 0
# /bin/cat >> /etc/stripetab
/dev/rstripe0 /dev/rxdv01h /dev/rxdv11h
CTRL D
# /etc/mkstripe -a
# /etc/makefs -t tfs /dev/rstripe0
# /bin/mkdir /projectx
# /etc/mount /dev/stripe0 /projectx
# /bin/cp /etc/fstab /etc/fstab.old
# /etc/mount -p > /etc/fstab
```

You can now create directories and files in the new file system. The directories and files appear logically under the directory that is the mount point and reside physically on the partition(s) holding the file system. If you add directories for users, you should change the directory's owner and group to that of the user. The following example adds a directory usr under /projectx, then adds a directory for hacker under /projectx/usr. (Use cat /etc/group|grep hacker to get the group name.)

```
# cd /projectx
# mkdir usr
# cd usr
# mkdir hacker
# chown hacker hacker
# chgrp algorithms hacker
```

The /etc/umount command unmounts a file system; for a striped file system, also execute /etc/unmkstripe. The data in the partition or disk stripe becomes unavailable. Any files that existed under the mount point in the root file system become available.

## 4.3 DUMPING, RESTORING, AND MOVING FILE SYSTEMS

The administration commands /etc/dump and /etc/restore are tools to protect against the loss of a file system (for example, a disk fails) and the loss of parts of the file system (for example, a user accidentally deletes a file). These commands can also be used to move file systems from one physical location to another; the general user commands bfio, tar, mv, and cp can also be used to move file systems and parts of file systems.

### 4.3.1 Dump Levels

Dumps can be taken at ten levels numbered 0 through 9. A level 0 dump backs up an entire file system. The other level dumps back up only those files modified since the last dump at a *lower* level. The use of levels permits you to balance the time required to dump file systems against the time (and trouble) required to restore files. As a starting point, consider the following 28-day cycle for dumping file systems:

**Dump Schedule**

| level | Mon | Tue | Wed | Thu | Fri | Sat | Sun |
|---|---|---|---|---|---|---|---|
| week 1 | 0 | 3 | 2 | 5 | 4 | 7 | 6 |
| week 2 | 1 | 3 | 2 | 5 | 4 | 7 | 6 |
| week 3 | 1 | 3 | 2 | 5 | 4 | 7 | 6 |
| week 4 | 1 | 3 | 2 | 5 | 4 | 7 | 6 |
| | 01 | 013 | 012 | 0125 | 0124 | 01247 | 01246 |

needed to reconstruct file system (level 1 not used in week 1)

A level 0 dump captures the entire file system but requires a significant amount of time. You should save monthly level 0 dumps indefinitely in a safe place (off-site is a good idea) so that you can recreate file systems and retrieve files that existed at any point in time (on a monthly basis). For a few small file systems, you might perform level 0 dumps weekly in place of the level 1 dumps.

The weekly level 1 dumps capture all changes since the monthly level 0 dump. Each week, the level 1 dump consumes more time because (typically) more files have changed since the beginning of the month. During the week, higher level dumps capture changes since either the day before or two days before, consuming smaller amounts of time than the level 1 dumps.

At any time, the dumps needed to restore a file system or retrieve any file are as follows: the level 0 dump; the most recent level 1 dump (except for week 1); the remaining dumps for the week, except where a lower numbered dump has been taken the next day. For example, to restore everything from the beginning of the month until Friday of week 3, you need the level 0 dump, the week 3 level 1 dump, the week 3 level 2 dump (Wednesday), and the week 3 level 4 dump (Friday). You do not need the week 3 level 3 dump (Tuesday) because the level 2 dump includes both Tuesday and Wednesday; likewise, you do not need the week 3 level 5 dump (Thursday).

Users should be aware of one weakness of the system: files with a short life span are likely to be lost. For example, if you create a file on Monday and accidentally delete it on Friday, the file is lost once the level 1 backup for the next week occurs (and the old dump tapes are recycled). Even for files of long duration, any modifications made during the week would be lost once the new level 1 backup occurs and the old tapes are recycled. Users must take care to recover lost files in the week in which the loss occurs.

### 4.3.2 Dump Procedure

Select a time of low use, notify your users, and shut down to single-user mode before dumping file systems. Dumping an active file system can produce inconsistencies in the dump and make restoration difficult. (Always shut down to single-user mode to dump the root partition. Administrators often take dumps on other partitions in multi-user mode due to the difficulty of scheduling shutdown time; however, this practice is not recommended.)

Issue an appropriate /etc/dump command for each file system; the next section explains the options concerning media. The terminal from which the command is issued becomes the recipient for many messages and queries (requiring responses) and, practically, is unavailable for other purposes until the dump is done. The n option causes all users in group operator to be notified when the dump needs attention (not really needed if you are in single-user mode). The u option logs the dump by adding a line to the file /etc/dumpdates and should be included for all dumps.

The following example performs a level 1 dump of the file system on /dev/rschs (the Exabyte cartridge drive). The u option is for logging the dump. The b option takes the corresponding option 32 to specify the blocking factor for write operations in 1K blocks. The d option takes the corresponding option 6250 to specify tape density. The s option takes the corresponding option 30000 to specify tape size. The f option takes the corresponding argument /dev/rschs to specify the output device.

```
# /etc/dump 1ubdsf 32 6250 30000 /dev/rschs /dev/rxdv01h
```

Time permitting, you should list the contents of the dump tape with the administration command /etc/restore. This procedure verifies that the dump is good and provides a listing for future reference. If the restore fails for any reason, take the dump again.

```
# /etc/restore tbf 32 /dev/rschs
```

Since the commands are long and you need one command for each file system, you should write a script to issue the appropriate commands for a dump of your system. You can write separate scripts for the separate days or insert logic into one script to dump at the correct level. The following example dumps a system containing five file systems. Logic in the script determines the dump level from the day of the week; the logic can be overridden (for example, for the monthly level 0 dump).

```
#! /bin/csh -f
#
# Shell script to take the correct level dump for the day of the week
# Can override by supplying the level as argument 1
#
if ( $1 == '' ) then
        set DATE = `date`
        set LEVELS = ( 1 3 2 5 4 7 6 )
        set DAYS = ( Mon Tue Wed Thu Fri Sat Sun )
        foreach i ( 1 2 3 4 5 6 7 )
                if ( $DAYS[$i] == $DATE[1] ) then
                        set n = $LEVELS[$i]
                        break
                endif
        end
else
        set n = $1
endif
/etc/dump $n'ubdsf 32 6250 30000 /dev/rschs /dev/rxdv00r'
/etc/dump $n'ubdsf 32 6250 30000 /dev/rschs /dev/rxdv00f'
/etc/dump $n'ubdsf 32 6250 30000 /dev/rschs /dev/rxdv01h'
/etc/dump $n'ubdsf 32 6250 30000 /dev/rschs /dev/rxdv10h'
/etc/dump $n'ubdsf 32 6250 30000 /dev/rschs /dev/rxdv11h'
```

To start a dump, mount a scratch cartridge or reel (unless the dump medium is other than tape) on the appropriate drive. Issue the command or shell script. As the dump proceeds, messages inform you of the approximate length of the dump and media requirements.

The dump stops and a query is issued if a tape volume fills. Replace the filled volume with another scratch volume, and type y on the terminal controlling the dump. You are notified when the dump terminates.

Ensure that each dump volume contains a physical label with at least the following information: date, dump level, volume number, partition name(s). Volume numbers are best written as "1 of *n*," "2 of *n*," and so on through "*n* of *n*." Include the blocking factor, cartridge specification, density, and size where the default is not used.

When backing up Diagnostix-2800, Concentrix-2800 must not be running. You can dump Diagnostix-2800 to any tape whose controller is attached to the VMEDP.

## 4.3.3 Dump Media

The typical dump medium is an Exabyte tape cartridge or a 9-track tape reel. You can dump to a file, space permitting, but files are not practical except for storing small dumps for a short time. Dumping to standard output in conjunction with restoring from standard input is a technique for moving file systems; see below.

The dump stops for a volume change based on the specified size and density of the volume, *not* when the volume reaches the end-of-tape mark. The following specifications use the standard tape volumes most efficiently (without running off the volume).

**Tape Size Specifications**

| | b option | c option | d option | s option |
|---|---|---|---|---|
| reel-to-reel 2400' | default | no | 6250 | default |
| Exabyte - 120 min. | 32 | no | 6250 | 30000 |
| Exabyte - 60 min. | 32 | no | 6250 | 15000 |
| Exabyte - 45 min. | 32 | no | 6250 | 11250 |

An Exabyte 120-minute tape holds 2GB and takes approximately 2 1/2 hours to fill.

A standard reel-to-reel 2400-foot tape volume holds approximately 150 megabytes and takes approximately 30 minutes to fill (at high density).

You can write multiple dumps to one tape by dumping all but the last file system to the no-rewind device:

```
#! /bin/csh
# dump all partitions to one tape
/etc/dump 2ubdsf 32 6250 30000 /dev/nrschs /dev/rxdv00r
/etc/dump 2ubdsf 32 6250 30000 /dev/nrschs /dev/rxdv00f
/etc/dump 2ubdsf 32 6250 30000 /dev/nrschs /dev/rxdv001h
/etc/dump 2ubdsf 32 6250 30000 /dev/nrschs /dev/rxdv0010h
/etc/dump 2ubdsf 32 6250 30000 /dev/nrschs /dev/rxdv0011h
```

You can use the e option of /etc/dump to calculate tape usage before taking the actual dump.

To access a dump that is not the first dump on a tape volume, use the s modifier of /etc/restore or the fsr option of mt.

### 4.3.4 Restoring a Full File System

To completely restore a file system, for example, if a disk fails and all data is lost, take the following steps:

1. Format the new or repaired disk, if necessary. See Format_disk in Appendix B.
2. Relabel the disk, if necessary. See Label_disk in Appendix B.
3. Make and mount the file system. See Make_file_system in Appendix B.
4. Set the current working directory to the mount point for the file system and restore the file system from the dumps. Use the r option of /etc/restore (see below).
5. Remove the file restoresymtable from the root directory of the restored file system.

Restore dumps in the following sequence:

6. Latest level 0 dump.
7. Latest level 1 dump following the level 0 dump, if any.
8. Latest level 2 dump following the previous dump, if any.
9. Latest level 3 dump following the previous dump, if any.
10. and so on...

For example, if you lose the disk after the Thursday dump during week 3 (of the prototype cycle shown above), restore the following dumps in the following order: level 0 dump, level 1 (Monday) dump for week 3, level 2 (Wednesday) dump for week 3, and level 5 (Thursday) dump for week 3.

For each dump, load the first volume containing the dump and type the /etc/restore command. (If you specified a blocking factor for the dump, specify the same blocking factor for the restore.) Specify the r command (full restore) and the f modifier (to name the input device), for example:

```
# /etc/restore rbf 32 /dev/rschs
```

You receive queries to change volumes – load the next volume and type y. You receive queries if the level is wrong for a volume – load the correct level volume and type y. You may receive a query concerning the ownership of the mount point – type y or n. If you took the dump on an active file system, you are likely to receive errors concerning missing and unexpected files – unless you have another dump, you have little recourse but to proceed by typing y.

After restoring a file system, take a fresh level 0 dump. Do not rely on the old dump for selective restoration of files because the inodes are different.

### 4.3.5 Restoring Files Selectively to a File System

If you know the exact path names of the files and directories that you want to restore, you can use the x option of /etc/restore to restore lost data. Otherwise, use interactive mode (i option).

To locate a recently lost file, work backwards through the dumps, from the last incremental dump taken to the last level 0 dump taken, skipping any dump at a higher level than one already searched. To locate lost files from some time ago, search backwards through the level 0 dumps.

Set the current working directory to the mount point for the file system. For each dump, load the first volume containing the dump and type the /etc/restore command. (If you specified a blocking factor for the dump, specify the same blocking factor for the restore.) Specify the x or i option (extract or interactive) and the f modifier (to name the input device). The following example restores a file:

```
# /etc/restore xbf 32 /dev/rschs /projectx/src/xyz.f
```

The following example restores a directory and all its contents (recursively):

```
# /etc/restore xbf 32 /dev/rschs /projectx/src
```

If you receive a query concerning the volume number, type the number of the volume that is loaded or load the next volume and type its number. If you receive a query concerning owner/mode, type y to set the ownership and permissions of the current working directory (the mount point for the file system that you are restoring) to what is recorded on the dump tape; type n to leave them.

In interactive mode, a simple command set permits you to examine the dump tape and restore files selectively. The prompt in interactive mode is as follows:

```
restore>
```

Commands can be abbreviated to one character. Type h or ? to display the available commands. See restore in Appendix A for a more complete description.

For multivolume sets, load the last tape first and work backwards.

The following example locates and restores a file in interactive mode:

```
# /etc/restore ibf 32 /dev/rschs
restore> cd files
restore> ls
 file1   file2   file3   file4   file5   file6   file7   file8
restore> add file1
restore> ls
*file1   file2   file3   file4   file5   file6   file7   file8
restore> extract
You have not read any tapes yet.
Unless you know which volume your file(s) are on you should start
with the last volume and work towards the first.
Specify the next volume #: 1
set owner/mode for /./? [yn] y
restore> quit
```

The add command adds files to the extraction list but does not actually restore the file system on disk. (The add command without an argument adds all files in the current working directory.) The extract command performs the actual restoration. Note that the ls command places an asterisk in front of each file that is on the extraction list.

### 4.3.6 Moving a File System

To move a complete file system or remake a file system (for example, to change its size):

1. Perform a level 0 dump of the file system using /etc/dump.
2. Make and mount the new file system. See Make_file_system in Appendix B.
3. Change directory to the new file system (cd command).
4. Restore the contents of the dump to the new file system using the r option of /etc/restore.

Dumps can be restored without concern for the original size of the file system. The file system can be completely or selectively restored.

If you can make a new file system before taking the dump, you can pipe the output of /etc/dump to /etc/restore and a dump file is not necessary. The following example moves the contents of the file system mounted at /py to /pz:

```
# /etc/dump 0f - /py | (cd /pz; /etc/restore xf -)
```

### 4.3.7 Saving and Moving Directories with bfio

The Concentrix-2800 bfio (file and archive manipulation) command saves and restores files and directory trees below the file system level. The bfio command is a general user command documented in the *Concentrix-2800 Commands and Applications Manual*. Use the -o option to archive files and the -i option to restore files.

The following command writes the directory tree starting at /px/admin to the tape volume loaded as /dev/rschs in bfio format:

```
# cd /px
# find admin -print | bfio -ov /dev/rschs
```

The following command restores under the /py directory the tape volume loaded as /dev/rschs. The tape must be in bfio format.

```
# cd /py
# bfio -iv /dev/rschs
```

Note that the contents of the bfio file are restored directly below the current working directory. Your current working directory should be one level above the top directory in the afio file, not the top directory itself. Directories and files that exist in the bfio file and the disk area being restored are overwritten.

You can move directory trees without creating a bfio file. The following example moves /px/admin to /py/admin:

```
# cd /px
# find admin -print | bfio -pv py
```

### 4.3.8 Saving and Moving Directories with Tar

The tar (tape archive) command saves and restores files and directory trees below the file system level. The tar command is a general user command documented in the *Concentrix-2800 Commands and Applications Manual*. Use the c command to write a tar file and the x command to read a tar file.

The following command writes the directory tree starting at /px/admin to the tape volume loaded as /dev/rschs in tar format:

```
# cd /px
# tar cf /dev/rschs admin
```

The following command restores under the /py directory the tape volume loaded as /dev/rschs. The tape must be in tar format.

```
# cd /py
# tar xpf /dev/rschs
```

Note that the contents of the tar file are restored directly below the current working directory. Your current working directory should be one level above the top directory in the tar file, not the top directory itself. Directories and files that exist in the tar file and the disk area being restored are overwritten.

The p modifier restores files with the same permissions as saved on the tar file (rather than basing permissions on the current file creation mask). File ownerships are as on the tar file.

Tar files are restricted to one volume.

You can pipe the output from one tar command to another tar command to move directory trees without the use of a tar file. The following example moves /px/admin to /py/admin:

```
# cd /px
# tar cf - admin | (cd /py; tar xf -)
```

The general user command dd is useful for reading media produced on other systems and for performing various conversions while copying data. See the *Concentrix-2800 Commands and Applications Manual*.

The general user command cp is not recommended for system administration functions because it does not preserve ownerships and modes.

## 4.3.9 Restoring The Root Partition

Listed below are the procedures for restoring a Diagnostix-2800 root partition, and a Concentrix-2800 root partition. You may need to use these instructions if a root partition is accidentally corrupted, or the disk is damaged and replaced.

### 4.3.9.1 Restoring the Diagnostix-2800 Root Partition

1. Insert the Concentrix-2800 distribution tape in the Exabyte drive.
2. Turn the power switch on the front panel to the ON (12:00) position, and press the RESET button.
3. Type o, B RETURN (as shown below) at the console prompts to start the system and stop in VMEbus booter mode:

```
cons>Options = B RETURN
```

4. Select the tape drive as the device containing the booter program: (you can type ? for help at any of the prompts in the steps below):

```
cons>device boot from
(d)isk (t)ape = t
```

   Note: You only need to type d and t. The system fills in the rest of the input.

5. Select the default controller for the tape drive by pressing RETURN:

```
controller (0-7) [return for default] = RETURN
```

6. The system loads the booter. After the message announcing the Alliant VMEbus Booter appears, type %r RETURN to direct the booter to copy files onto the RAM disk, as shown below:

```
Alliant VMEbus Booter [made: Wed May 9 19:32:23 EDT 1990]
Type %h for help

boot from:  xdv(0,0,p)dunix: %r RETURN
```

7. Direct the booter to load the RAM disk from the second file on the tape, as shown below:

```
load ramdisk from:  xdv(0,0,p)ramdisk.init: hs(2)  RETURN
```

8. Direct the booter to load Diagnostix-2800 from the next file on the tape, as shown below:

```
boot from:  hs(2)dunix: hs(*) RETURN
```

9. Direct the booter to load the dddt program from the next file on the tape, as shown below:

```
boot from:  hs(*)dddt:  hs(*)RETURN
```

10. When the Diagnostix-2800 Install Menu appears, select option I, and press RETURN:

**Diagnostix-2800 Install Menu**

```
                    Diagnostix Install Menu

        (Please Choose one of the following options and hit <return>.)

Option  Action                                              Default Command
------- ------------------------------------------------- ------------------
   I    Create root file system and Install Diagnostix     (/bin/iu.dx.sh i)
        (includes loading hard disk booter).
   U    Upgrade Diagnostix operating system.               (/bin/iu.dx.sh u)
   S    Boot to Diagnostix.                                (/etc/reboot -S)
   H    Boot to cons> prompt.                              (/etc/reboot -H)
   X    Exit to Diagnostix prompt.

Option? > I RETURN
```

11. Select the default root disk for Diagnostix-2800:

```
>> Your selection? 1 RETURN
```

The system will update Diagnostix-2800. This process takes about 5 minutes.

12. From the Diagnostix-2800 Install Menu, choose option (upper case) S to reboot, stopping in Diagnostix-2800:

```
Option? > S RETURN
```

13. The reboot occurs. When you get the Diagnostix prompt, type the following command:

```
dunix-sh# dist/dx_install I RETURN
```

Ignore any messages that are followed by "Exit 1 (ignored)."

14. If you have a backup tape you wish to restore, remove the installation tape and put in your backup tape. To restore the entire root partition, enter the following command. When the restore is finished (or if you are skipping the restore), reboot the system:

```
dunix-sh# restore rbf 32 /dev/rschs /dev/rxdv00p

dunix-sh# /etc/reboot
```

### 4.3.9.2 Restoring the Concentrix-2800 Root Partition

1. Turn the power switch on the front panel to the ON (12:00) position, and press the RESET button.

2. Type o, S RETURN, u (as shown below) at the console prompt to start the system and stop in the Diagnostix-2800 shell prompt :

```
cons>Options =  Enter RETURN to keep, SPACE to clear, or new op-
tions.
New Options? S RETURN
cons>u
```

3. Enter CTRL d at the dunix-sh# prompt.

```
dunix-sh# CTRL d
```

   The following messages appear:

```
FX/2800 Startup

Entering Console Boot Loop.
```

4. At the series of questions about reboot with flags, enter R to access the RAM disk, e to enable booting from the RAM disk, and s to stop in single-user Concentrix-2800, followed by RETURN each time, as shown below:

```
Reboot with flags set to " "? (y for yes, ? for help): R RETURN

Ramdisk booting is currently disabled.
Do you want to enable (e) or disable (d)? e RETURN
Reboot with flags set to " "? (y for yes, ? for help): s RETURN
```

5. When the Concentrix-2800 Utilities Menu appears, select option I, and press RETURN:

**Concentrix Utilities Menu**

```
                    Concentrix Utilities Menu

        (Please Choose one of the following options and hit <return>.)

 Option  Action                                          Default Command
 ------- ----------------------------------------------- -------------------
    I    Create root (/) and /usr file systems and
         Install Concentrix                              (/bin/iu.cx.sh i)
    U    Upgrade Concentrix operating system             (/bin/iu.cx.sh u)
    NR   Nonstandard upgrade of Concentrix root (/)
         partition only                                  (/bin/iu.cx.sh nr)
    NU   Nonstandard upgrade of Concentrix /usr
         partition only                                  (/bin/iu.cx.sh nu)
    R    Boot to Multi-User Concentrix                   (/etc/reboot)
    s    Boot to Single-User Concentrix                  (/etc/reboot -s)
    S    Boot to Diagnostix                              (/etc/reboot -S)
    H    Boot to cons> prompt                            (/etc/reboot -H)
    X    Exit to Concentrix prompt.

Option? > I RETURN
```

6. Select the default root disk for Concentrix-2800:

```
>> Your selection? 1 RETURN
```

The system will update Concentrix-2800. This process takes about 20 minutes.

7. From the Concentrix-2800 Install Menu, choose option (lower case) s to reboot, stopping in single-user Concentrix-2800:

```
Option? > s RETURN
```

8. The reboot occurs. When the prompt appears, type the following command:

```
Concentrix-sh# dist/cx_install I RETURN
```

Ignore any messages that are followed by "Exit 1 (ignored)."

9. If you have a backup tape you wish to restore, put your backup tape in the drive. To restore the entire root and usr partitions (in the example, root is the first file on the tape and usr is second), enter the following commands:

```
Concentrix-sh# restore rbf 32 /dev/rschs /dev/nrxdv00r
Concentrix-sh# restore rbf 32 /dev/rschs /dev/rxdv00f
```

10. When the restore is finished (or if you are not doing a restore), reboot the system:

```
Concentrix-sh# /etc/reboot
```

## 4.3.10 Cleaning Up /tmp

When you delete files, Concentrix-2800 does not collapse the depleted directories. Although the script /etc/rc deletes /tmp files at system startup, the directory itself is left intact, with all its null entries. To assure faster execution of programs that use /tmp files, you can change the /etc/rc script to delete the directory (not just the files in it). See the example below:

```
#!/bin/sh

tmp="/tmp"                      # In case we want to use this elsewhere.
new=".rc.new"
old=".rc.old"
preserve="lost+found quotas"    # Files and/or directories to preserve.

if [ -h "$tmp" ]
then                            # This is the symbolic link case.
    if [ ! -d $tmp ]
    then                        # pointed-to directory does not exist.
        echo "The symbolic link to where $tmp points to does not exist."
        tmp="/tmp.is.non.existent"
        rm -rf $tmp
        mkdir $tmp
```

```
    else                            # pointed-to directory exists.
        cd $tmp
        tmp="`/bin/pwd`"
        if [ -z "$tmp" -o "$tmp" = "/" ]
        then                        # evil directory (e.g. /tmp -> /).
            echo "Invalid temporary directory symbolic link to $tmp"
            tmp="/tmp.is.invalid"
            rm -rf $tmp
            mkdir $tmp
        fi
    fi

cd /

elif [ ! -d "$tmp" ]
then                                # /tmp does not exist at all.
    echo "There is no $tmp directory. We will make one for you."
    rm -rf $tmp
    mkdir $tmp
fi

tmpnew="$tmp$new"
tmpold="$tmp$old"

rm -rf $tmpnew $tmpold
mkdir $tmpnew
chmod a+rw $tmpnew
if [ -n "$preserve" ]
then
    cd $tmp
    mv -f $preserve $tmpnew
    cd /
fi

mv $tmp $tmpold
mv $tmpnew $tmp
rm -rf $tmpold &
```

## 4.4 TUNING AND FIXING

The /etc/tunefs command and various arguments to the /etc/makefs command permit a more detailed specification of a file system. Typically, the guidelines provided by Make_file_system in Appendix B suffice and the detailed information presented in this section is not needed.

The command /etc/fsck fixes file systems left in an inconsistent state by an improper shutdown or a hardware failure. Typically, the command executes in automatic mode during system start-up and manual intervention is not needed. The detail in this section provides some background for manual intervention; also see fsck in Appendix A.

### 4.4.1 Examining a File System

The command /etc/dumpfs provides information on a file system. The first part of the display is based on the file system's superblock; the remaining parts of the display provide specific information concerning each cylinder group in the file system. As you become familiar with the structure of a file system, the abbreviations and data presented by /etc/dumpfs will make sense.

The following example dumps the file system information for a file system made on the standard h partition of an fj690s disk (Fujitsu M2344) and shows just the superblock (first) part of the display. The magic number 11954 identifies the file system as type 4.3 (the magic number 81486 identifies a track-allocated file system).

```
# /etc/dumpfs /dev/rxdv01h

magic   11954   time    Mon Mar 26 15:15:12 1990
nbfree  95395   ndir    392     nifree  82379   nffree  844
ncg     47      ncyl    742     size    801360  blocks  788939
bsize   4096    shift   12      mask    0xfffff000
fsize   1024    shift   10      mask    0xfffffc00
frag    4       shift   2       fsbtodb 1
cpg     16      bpg     4320    fpg     17280   ipg     2048
minfree 10%     optim   time    maxcontig 1     maxbpg  256
rotdelay 4ms    headswitch 0us  trackseek 0us   rps     60
ntrak   27      nsect   80      npsect  82      spc     2160
trackskew 0     interleave 1
nindir  1024    inopb   32      nspf    2       id      0x000000000
sblkno  12      cblkno  16      iblkno  20      dblkno  276
sbsize  2048    cgsize  4096    cgoffset 40     cgmask  0xfffffffe0
csaddr  276     cssize  1024    shift   8       mask    0xffffff00
cgrotor 40      fmod    0       ronly   0
```

## 4.4.2 File System Checks

The superblock (see below) and other structures within the file system reflect the current state of a file system. Changes to the status information, as files are updated, occur in memory buffers and are only flushed to disk as the buffers are reallocated. At any random time on an active file system, therefore, the status information on disk is not up to date.

Synchronization of the file system means that all file system information is flushed to disk to bring the file system up to date on disk. Synchronization of the file system occurs at the following times:

- Normal shutdown of the system.
- Every 30 seconds if /etc/update is executing; /etc/update is started in the standard /etc/rc script at system initialization.
- If the command /bin/sync (or the sync system call) is explicitly executed.

In addition, commands that manipulate file systems, such as /etc/mount, synchronize the file system.

Stopping the system without first synchronizing the file system leaves the file system in an inconsistent state. You should never deliberately leave the file system in an inconsistent state, for example, by turning off the power while the system is running; always follow standard shutdown procedures (see Chapter 1).

Typically, the file system is left in an inconsistent state due to hardware and software failures. If a hardware or software failure stops the system, the file system must be checked and repaired when the system starts again. The command /etc/fsck checks and repairs file systems in two modes:

- **Non-interactive** – The –p option indicates non-interactive mode, in which repairs are automatically applied to a subset of the possible inconsistencies. Unexpected inconsistencies cause /etc/fsck –p to exit with an error status.
- **Interactive** – No options or options other than –p indicates interactive mode, in which you are prompted for concurrence before repairs are applied.

The command /etc/fsck –p is executed at system initialization from the standard /etc/rc script unless a fastboot is in effect. Do not fastboot if the system stopped due to a failure or you suspect a problem with the file system.

Interactive mode is necessary if /etc/fsck –p exits with an error status or if you suspect a problem in a file system (for example, if you receive a hard disk error message). See fsck in Appendix A for detailed information.

## 4.4.3 File System Layout

The following table presents the performance parameters that can be affected by /etc/makefs and /etc/tunefs. A more detailed discussion follows in the subsections.

**File System Performance Parameters**

| | defaults | | command options | |
|---|---|---|---|---|
| description | 4.3BSD | track | makefs | tunefs |
| block size | 4096 | 1 track[1] | -b[2] | none |
| blocks per transfer | 1024[5] | n/a | -a | -a |
| cylinders per group | 16 | 24 | -c | none |
| file limit in cylinder | 20%[3] | n/a | -e | -e |
| fragment size | 1024 | 4096 | -f | none |
| full blocks | n/a | 1 | -F | none |
| indirect pointer block | n/a | fragsize | -i | none |
| bytes per inode | 2048[4] | 2048[4] | -I | none |
| minimum free space | 10% | 10% | -m | -m |
| rotational delay | 5[5] | n/a | -d | -d |

[1]Preferably. The block size must be a multiple of 4096; the number of blocks per cylinder must be even.
[2]Track-allocated file systems only; 4.3BSD file systems must be 4096.
[3]Approximately.
[4]Or the fragment size if it is less than 2048.
[5]Changing this value with tunefs may not have the desired effect if substantial space is used on the disk, since the pattern of free space may reflect parameters for filesystems created pre-Concentrix-2800 2.2.

### 4.4.3.1 Cylinder Groups

A file system is divided into cylinder groups. Each cylinder group consists of an equal number of consecutive full cylinders, except that the last cylinder group consists of the remainder of the cylinders. By default, the size of a cylinder group is 16 cylinders for a 4.3BSD file system and 24 cylinders for a track-allocated file system. You can explicitly specify the cylinders per group with the -c option to /etc/makefs.

The purpose of cylinder groups is to cluster and disperse data for optimal access to the file system. If at all possible, the files in one directory (except for directory files) are placed in the same cylinder group for speed of access. A new directory, on the other hand, is placed in the cylinder group with the most free space to avoid overflows in the cylinder groups. If at all possible, the data for one file is placed rotationally optimal on one cylinder, or within one cylinder group if one cylinder is impossible. However, large files are continued in a new cylinder group at 32 kilobytes and at every megabyte thereafter to avoid overflows in the cylinder groups.

A UFS file system with a few large files benefits by raising the amount of the file allowed in one cylinder group. You can adjust the limit with the -e option of /etc/tunefs (which can be applied to an existing file system). Note that this option does not apply to TFS filesystems. For TFS, the system switches cylinder groups only when running out of space or when starting a new indirect block.

The following example illustrates division of a 4.3BSD file system on the standard r partition of an fj1000s disk (Fujitsu M2382) into cylinder groups. The disk contains 2160 512-byte sectors per cylinder.

**Cylinder Groups for File System in Standard r Partition**

| cg number | cylinders | size in bytes |
|---|---|---|
| 0 | 16 | 17,694,720 |
| 1 | 16 | 17,694,720 |
| 2 | 4 | 4,423,680 |
| totals | 36 | 39,813,120 |

### 4.4.3.2 Inodes, Blocks, and Fragments

A file system contains the following structures:

- **Inode** – An inode describes a file and points to the data (for a data file). The descriptive information includes: user and group ownership, type (data, directory, character special, block special, symbolic link, or socket), permissions (read, write, execute), time stamps (file last modified, file last accessed, inode last modified), number of links to file, and file size. Data pointers 0–11 point directly to data blocks; data pointer 12 points to data blocks through an indirect block; data pointer 13 points to data blocks through two levels of indirect blocks; and data pointer 14 points to data blocks through three levels of indirect blocks.

- **Data blocks** – Data is stored in blocks of a fixed size. The block size of a 4.3BSD file system is 4096 bytes; the default block size of a track-allocated file system is the track size (or multiple tracks if necessary to make the blocksize a multiple of 4096) and can be adjusted with the -b option of /etc/makefs. In a 4.3BSD file system, the twelve direct pointers per inode provide 49,152 bytes of direct storage.

- **Indirect data blocks** – If a file size exceeds direct storage requirements, data pointer 12 points to a block that contains more pointers; these pointers point to the data blocks. By default, an indirect data block is the same size as a data block; for TFS filesystems, the size can be adjusted with the -i option of /etc/makefs. In a 4.3BSD file system with the default indirect block size, the 1024 indirect pointers provide an additional 4,194,304 bytes of storage. Data pointers 13 and 14 provide further direction for very large files.

Inodes are statically allocated when a file system is made. By default, one 128-byte inode is allocated for every 2048 bytes of file system space. You can modify the allocation scheme with the -I option of /etc/makefs.

The following example illustrates the structure of a file. The file contains 15 data blocks: the first twelve are direct; the last three are accessed through the indirect block to which pointer 12 points.

**Structure of a file**

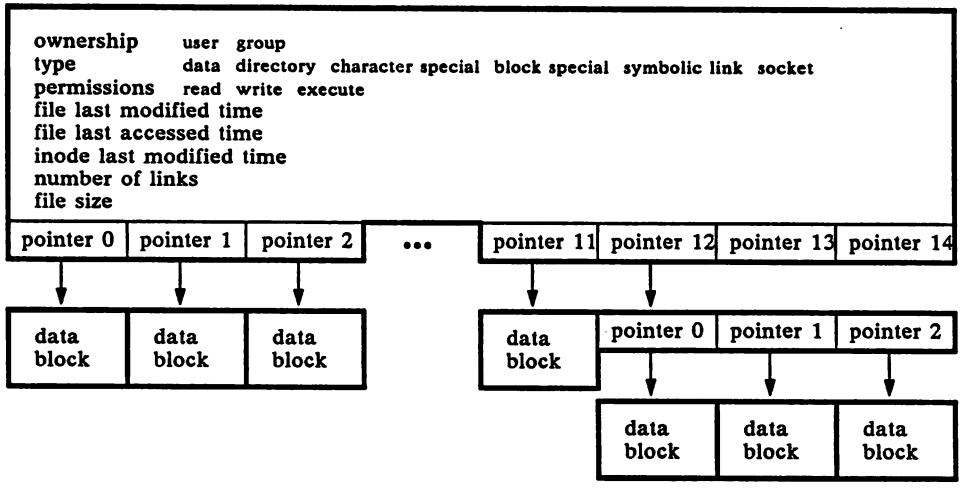


In directory files, the data blocks are specially formatted to contain the names and inode numbers of the files in the directory.

The use of large blocks permits the more rapid transfer of data between disk and memory, but wastes space when the file system contains many small files. Therefore, blocks are divided into fragments and fragments are allocated where whole blocks are not needed. By default, the fragment size is 1024 bytes on a 4.3BSD file system and 4096 bytes on a track-allocated file system. You can adjust the fragment size with the -f option of /etc/makefs.

A BSD4.3 file system block can consist of 1, 2, 4, or 8 fragments. The fragment size of a track-allocated file system must be a power of 2 multiple of 4096 bytes: 4K, 8K, 16K, 32K, and so on. The block size of the track-allocated system must be 1 to 32 times the fragment size, and the number of blocks per cylinder must be an integer.

For example, a 5012-byte file on a default 4.3BSD file system is stored as one 4096-byte block and one 1024-byte fragment. If the file decreases in size below 4096, the fragment is freed. If the file increases in size beyond 5120 bytes, another fragment is added.

### 4.4.3.3 Free Space

As the file system nears capacity, the ability to allocate space efficiently diminishes. By default, 10% of the file system is reserved as the minimum free space. If the threshhold is crossed, only the superuser can allocate further space in the file system. You can adjust the minimum free space with the -m option to /etc/makefs and /etc/tunefs (which can be applied to an existing file system).

### 4.4.3.4 Rotational Delay

In UFS filesystems made under Concentrix-2800 2.2, multiple consecutive blocks are allocated whenever possible. In this case, rotational delay controls the selection of a block on the same cylinder where a consecutive block is unavailable. The system attempts to get a block which can be accessed without loss of a revolution. The default value of 5 milliseconds assures this for all supported devices.

For other devices, you can determine an appropriate value by doing a tunefs -a 1 to prevent sequential allocation from obscuring the results, and then adjusting the delay and reading and writing files at different tunefs -d values, picking the value that gives the best performance. Remember to set tunefs -a back to 1024 afterwards.

### 4.4.3.5 Superblock

A file system contains a structure called the superblock. The superblock contains the basic parameters of the file system, such as the number of data blocks and a count of the maximum number of files. Because the superblock contains critical data, it is replicated through the file system on different cylinders and platters. When the file system is created, /etc/makefs displays the block numbers containing the replicated (alternate) superblocks. Write down and save these block numbers.

Associated with the superblock is non-replicated summary information. The summary information changes as the file system is modified. The summary information contains the number of blocks, fragments, inodes, and directories in the file system.

Each cylinder group includes a block map describing available blocks in the cylinder group, and summary information describing the usage of data blocks within the cylinder group. The cylinder group bookkeeping information begins at a floating offset from the beginning of the cylinder group. The offset for the *i+1*st cylinder group is about one track further from the beginning of the cylinder group than it was for the *i*th cylinder group. In this way, the redundant information spirals down into the pack; any single track, cylinder, or platter can be lost without losing all copies of the superblocks. Except for the first cylinder group, the space between the beginning of the cylinder group and the beginning of the cylinder group information stores data.

If the primary superblock is bad, specify a replicated (alternate) superblock to /etc/fsck to perform a file system check on a specified file system. Block 24 is always a replicated superblock (block 32 on older file systems); however, best practice is to note the replicated superblocks for future reference when you make the file system. For example, if the file system on /dev/rxdv01h has a bad superblock, perform a file system check as follows:

```
# /etc/fsck -b 24 /dev/rxdv01h
```

If a replicated superblock is incorrect, that is, does not agree with the primary superblock, an interactive file system check usually succeeds but a noninteractive file system check fails. You can force agreement by propagating the contents of the primary superblock to the replicated superblocks with the -A option of /etc/tunefs. You must also specify another option to force /etc/tunefs to work, for example, the default value for an option or the known tuned value. First perform an interactive file system check to be sure that the file system is correct. The following example fixes the replicated superblocks on a file system (10 is the default value for -m):

```
# /etc/fsck /dev/rxdv01h
# /etc/tunefs -A -m 10 /dev/rxdv01h
```

# CHAPTER 5
# MONITORING AND TUNING

### Administration Files

| name | description |
|---|---|
| /console/console.config | console configuration file in Diagnostix-2800 root partition (edit) |
| /dev/kmem | kernel memory ** |
| /etc/rc.cluster | system startup cluster configuration (edit) |
| /etc/usrclass | default scheduling classes file ** |
| /unix | system kernel |
| /usr/cs/elog | hardware error log |
| /usr/adm/messages | system messages file |
| /usr/lib/crontab | clock execution file (edit) |

** See *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

### Administration Commands (Appendix A)

| name | description |
|---|---|
| clcr | create cluster |
| clinfo | print cluster information |
| clrm | remove cluster |
| cron | clock daemon |
| dmesg | collect system diagnostic messages in error log |
| setclass | set process class |
| setclassdef | change classdef flag in kernel image |
| setcres | set computing resource |
| setpn | set the maximum number of processors for a concurrent job |
| settype | set computing resource type |
| setsched | set scheduling vector for computing resource |
| showsched | display scheduling vectors |
| vmlock | enable/disable process swapping and outpaging |

### Administration Procedures (Appendix B)

| name | description |
|---|---|
| Build_kernel | build custom kernel |
| Patch_kernel | modify variables in the system kernel |

**General Commands (Concentrix-2800 Commands and Applications Manual)**

| name | description |
|---|---|
| adb | assembler-level debugger |
| df | report disk free space |
| du | summarize disk usage |
| execute | run program on specific processor type |
| find | search directory tree for specific files |
| grep | search a file for a pattern |
| kconfig | display system memory usage, process statistics, or mbuf usage |
| kill | terminate job or process |
| ls | list contents of directory |
| mon | monitor process/system status (crt display) |
| ps | display process status information |
| syscfg | show system configuration |
| w | show who is on and what they are doing |
| xproc | display processes or process times (user/system, IP, CE and COMPLEX), or activity on the individual processors. |

## 5.1 ACTIVE SYSTEM

Typically, performance degradation is due to too many processes, high resource usage by one or more processes, or some combination of these factors. You can alleviate performance problems by using a batch queue for large jobs and restricting the number of batch devices assigned to the queue, avoiding jobs that burden the file system (such as long copies and finds, especially on a system using NFS) during prime time, reducing the number of interactive users logged on at one time, and increasing physical memory and disk space.

### 5.1.1 System Configuration and Current Users

The command /usr/alliant/syscfg provides information on the configuration of your system, including the amount of memory, the number of CEs, the number of IPs, and the sizes of the paging and swapping files. Note the following example:

```
# /usr/alliant/syscfg
System Configuration for FX2800 buddy (serial # 0 ):
    Memory Size    = 256 Mb
    CE-Cache Size = 2048 Kb
    User Memory    = 224.7 Mb
    Paging File    = 54.0 Mb
    Swapping File = 216.0 Mb
    Number of IPs = 1
    Number of CEs = 12
       Detach/Attach = 8/4
```

The –d option provides information on the devices connected to each IOM.

The command /usr/ucb/w displays information about current users, including login terminal names, remote system logged in from, time on the system, and the commands being executed:

```
# w

  2:54pm  up 19:04,  60 users,  load average: 1.75, 1.88, 1.65, 1.31
User     tty from            login@  idle   JCPU   PCPU  what

harald   p8  hugin:0.0       7:32am     4     13      1  csh
steve    p9  cygnet          7:34am  3:13   1:35      6  a.out
mat      pa  myzoo           7:37am    11  15:57      2  -csh  (csh)
mat      pb  myzoo           7:38am  6:46   5:14   5:14  mon -r -t
jim      q0  boris           2:17pm         2:58      2  adeno -xout
jack     qb  jackal          9:11am    40  10:04     17  vi qgetqueue.c
jack     qc  jackal          9:11am        43:24     59  vi struct.c
jack     qd  jackal          9:11am  2:14   5:53     12  vi fxenv
jack     qe  jackal          9:11am  2:20   2:16     12  ../daemon/interactor_d
fred     qf  fornax:0.0      9:12am     6   4:27   1:04  emacs
doug     r3  moosewood:0.0   9:21am    49     12      1  csh
doug     r4  moosewood:0.0   9:21am76days      2      1  csh
kevin    rb  babel           11:41am     3   5:30     29  emacs
bob      rc  cl2             11:25am           7      1  w
adele    re  kahala:0.0      11:49am          18      6  vi px_hilight.c
mark     s2  babel           2:38pm      2  10:50   3:53  /bin/csh runall O_mark
   ...
```

## 5.1.2 Monitor Display

The command /usr/alliant/mon displays resource usage information and can help locate problem processes. Useful invocations of the command are as follows:

- **mon -c** – Displays the state of each processor and the computational complex (if active): time doing real work, system time, user time, interrupt time, and lockwait time. The following example is for a system with 1 IP and 8 CEs, 4 of which form a cluster (CL).

```
          USER       SYSTEM ++      INTERRUPT --      LOCKWAIT **

1    IP    0    0%
.3   CLx4  0    0%
.7   CE    0    9%
.7   CE    1   22%
1    CE    2   17%
1    CE    3   20%
.7   CE    4   18%
.7   CE    5   17%
1    CE    6   30%
1    CE    7   47%
```

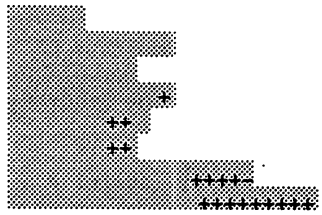

- **mon -t** – Displays the top cpu users, factoring in the relative speeds of the different computational types. The display includes the user name, process identification, percentage of cpu resource (IP, CE, or cluster), expended cpu time, physical memory size, and command being executed. Note the following example:

```
USER      PID   %IP     %CE  %CL CPUTIME  RSZ  COMMAND

foo      4500          10.8          0  152K  nethack
tony     7082           9.2       17:51  236K  mon -r -t
root     4533           9.0          4  220K  mon -t
root      446           5.3         10  100K  vi fxada.install
root      210           5.1          4  100K  /etc/rwhod
john     4474           1.8          3  128K  ls
root     1040           1.6       1:50   84K  in.rlogind 7d0000f3.1020
rich      204           1.6       1:07  140K  rn
root       57                    28:15  116K  /etc/ypserv
root        7           0.1       5:17  108K  /etc/routed
```

- **mon -m** – Displays the top users of real memory. The display includes the process identification, scheduling resource class, number of times swapped, virtual memory size, physical memory size, percentage of cpu resource (IP, CE, or cluster), page faults, page reclaims, and command being executed. Note the following example:

```
 PID  C  SW S   VSZ   RSZ   PF    PR STA USER    COMMAND

1907  3   0    2.1M  1.2M    0     0 Run stan    emacs vmul.207
5580  3   0    940K  316K    0     0 Stp joan    emacs summary_fortran
6862  2   0    2.8M  240K    0     0 Run root    fortran1 -nolink -noasso
 130  3   0    440K  232K    0     0 Slp root    /usr/lib/nqs/nqsdaemon
9553  3   0    364K  228K    0     0 Slp tont    mon -r -t -f user pid %i
```

A mon display is continuous, changing every so many seconds to reflect changing events on the system. You can modify the display by entering options interactively. For example, if you invoke mon for the top users display:

```
# /usr/alliant/mon -t
```

You can change to the physical memory display by typing m.

Type q to quit.

The /usr/alliant/kconfig command gives a continuous display of system memory usage (swap, page and global), process statistics, or mbuf usage. Unless the –i option is used, the display is updated every 5 seconds. The version number of the Concentrix-2800 release and information about shared memory, daylight saving time, remote port, and NFS installation kernel variables are also displayed.

Typing one of the option characters at any point switches to the appropriate display when kconfig is in continuous display mode. The time interval can be changed from 1 to 10 seconds by typing the appropriate character from 1 to 0 (0=10).

The following example shows the display for the kconfig –p option:

```
% kconfig -p

   Alliant Concentrix 2800 2.1.0 on 'marley'       Version# 20000003
   Built Tue Dec 18 12:09:24 EST 1990              Thu Jan 17 16:37:12 1991

context switches ..............  1305   number of swapped runnable procs     0
idle switches .................     0   super. mode ce based exceptions     99
number of sleeps causing switch   214   user mode ce based exceptions ..  2057
calls to trap .................     0   collapse complex events ........     0
calls to syscall() ............ 11003   expand complex events ..........     0
asynchronous system traps ......  257   dynamic cluster explodes .......     4
forks executed .................   55   dynamic reconfigs to cluster ...     4
obreak system calls ............  730   device interrupts .............. 1909
exec system calls ..............   52   pseudo-dma interrupts ..........     0
exec system calls (command file)    0   kfe control block requests .....    12
known file searches ............    0   kfe control block hits .........     0
known file hits ................    0

Concentrix will switch to daylight saving time
Only ROOT can use setclass
Remote console port    = DISABLED
NFS                    = ENABLED
```

## 5.1.3 Process Status

The command /bin/ps displays information about current users. Use the options aux for a user oriented display of all processes. The display includes login name, process identification, virtual memory size, physical memory size, terminal name, status (running, stopped, sleeping, idle), time on the system, and the command being executed.

```
# /bin/ps aux | /bin/more
USER       PID   VSZ     RSS TT STAT    TIME  COMMAND
root         0     4       0 ?  Slp    17:25  Swapper
root         1    48 L    44 ?  Slp    12:04  /etc/init
root         2     4       0 ?  Slp     0:00  Page reclaimer
root         3     4       0 ?  Slp     0:00  Global page reclaimer
root         4     4       0 ?  Slp     0:00  Primary page cleaner
root         5     4       0 ?  Slp    35:00  Diag Handler
root         6     4       0 ?  Slp     0:20  Printf daemon
roger      303   292 L   180 q2 Slp     0:12  -csh  (csh)
roger      322   104 L    68 q2 Slp     0:22  month -d
kathy      336   280 L   168 16 Slp     0:09  -csh  (csh)
paul       352   316 L   212 pd Slp     0:47  rn
root       464   328 L   204 13 Slp     0:04  make ALTROOT=/rev5src

...
```

To examine the processes that belong to one user, pipe the output of ps aux to grep *username*:

```
# /bin/ps aux | /bin/grep reid
reid       711   296 L   188 p0 Slp     0:06  -csh  (csh)
reid      1328   984 L   500 p0 Slp     0:30  dbx cos
reid      1372   192 L   128 p0 Stop    0:51  cos
root      4056    64 L    68 22 Slp     0:00  grep reid
```

ps also displays thread identifiers.The following example demonstrates the execution of a program (named watch) that creates two threads. While the program is running in the background, the ps command displays the threads. The PID field for each thread displays the same process identifier. For the root thread, the COMMAND field displays the name of the command that started the process; for the other threads, the COMMAND field displays the thread identifier. See the administration command thr_overview for more information about threads.

```
% watch&
func: this is thread 1
Thread id of thread 1: 28121
Thread id of thread 2: 3934
func: this is thread 2
[1] 32200
% ps

  PID TT STAT    TIME  COMMAND
19674 ra Run     0:00  ps
23904 ra Slp     0:00  -csh  (csh)
32200 ra Slp     0:31  watch
32200 ra Slp     0:00  [03934]
32200 ra Slp     0:00  [28121]
```

The /usr/alliant/xproc command gives a continuous display of processes (similar to ps) or process times (user/system, IP, CE and COMPLEX), or activity on the individual processors. Unless the –i option is used, the display is updated every second.

The various headings have the following meanings:

USER The user name of the process owner.
PID The process ID number.
STATE The state the process is in ... one of SLEEP, WAIT, RUN, IDLE, ZOMBY, or STOP.
RSZ The resident set size of the process.
WSZ The working set size of the process.
PRI The process priority.
NI The process 'nice' level.
CL The current class the process is running at.
PR The run queue index that the process is at.
CPU Which processor the process is/can run on.
FLAGS Memory management flags for the process ... one or more of SYS, NOMEM, SWAPPED, LOCK, NOSWAP and NOPAGE.
IP,CE,CMPLX
Which resource the accumulated time is displayed for. This is followed by the number of processors the process is running on (in () ), if applicable.
CLASS LIST
The list of classes (and their priority) the CPU will run. These are the classes defined by setsched.

```
% xproc -i9

PROC slots .. Used = 174   Free = 358          Elapsed Time = 0.00s
Files ..... Open = 440   Free = 1550
COMMAND            PID USER      STATE    RSZ    WSZ PRI NI CL PR CPU  FLAGS
cc -c -Og -I../. 27795 ama       STOP     56K   128K  30 0  3  0 IpCe
fxc1 -Og -I../.. 27855 ama       STOP    212K  4.5M  30 0  3  1 IpCe
sh -ce for i in  27399 eaw       SLEEP    32K   128K  30 0  3  0 IpCe
make JFLAG=-j5   27997 eaw       SLEEP   184K   308K  30 0  3  0 IpCe
rlogind           6255 root      SLEEP   160K   256K  26 0  3  0 IpCe
-csh             29164 hossein   SLEEP   200K   432K  40 0  3  0 IpCe
tail -f meme     29522 michaud   SLEEP    32K    68K  40 0  3  0 IpCe
rlogin flash     26995 shtil     SLEEP   112K   116K  28 0  3  0 IpCe
rlogin flash     31238 shtil     SLEEP   112K   116K  26 0  3  0 IpCe
vi Motif.tmpl    31615 ama       SLEEP   140K   448K  28 0  3  0 IpCe
vi site.def      31888 ama       SLEEP   144K   448K  28 0  3  0 IpCe
rlogind          30847 root      SLEEP   160K   256K  26 0  3  0 IpCe
-csh             31450 perron    SLEEP   124K   380K  40 0  3  0 IpCe
csh ./maketape.s 32181 hossein   SLEEP   192K   428K  40 0  3  0 IpCe
tar cvf /dev/rsc   115 hossein   SLEEP   100K   112K  20 0  3  1 IpCe
```

## 5.1.4 Process Killing (Hung Process)

A user process may hang in a state that locks the user's terminal. For example, a program under development that does not accept CTRL C interrupts might enter an infinite loop. First, ensure that the terminal really cannot be accessed: type CTRL Q to make sure the terminal is not simply stopped (or press the STOP key depending on the terminal type); type CTRL C to make sure the program cannot really be terminated; type CTRL Z to make sure the program cannot be suspended.

If you can suspend the process from the user's terminal with CTRL Z, type jobs to verify the job number of the suspended job, then kill the job. For example:

```
CTRL Z
Stopped
% jobs
[1]  + Stopped          cos
% kill %1
```

You might also try physically resetting the terminal or power cycling the terminal, especially if the display seems skewed.

If you cannot recover the user's terminal from the terminal, you must kill the offending job from another terminal. You may have to go so far as to kill the user's login process. On another terminal, login as root or the affected user and examine the user's current processes. Kill the offending process with a KILL (9) signal and verify the kill with another ps command. The following example kills the program cos for user reid:

```
# /bin/ps aux | /bin/grep reid
reid       711   296 L   188 p0 Slp     0:06  -csh  (csh)
reid      1372   192 L   128 p0 Run     0:51  cos
root      4056    64 L    68 22 Slp     0:00  grep reid
# /bin/kill -9 1372
# /bin/ps aux | /bin/grep reid
reid       711   296 L   188 p0 Slp     0:06  -csh  (csh)
root      4056    64 L    68 22 Slp     0:00  grep reid
```

If the user's terminal is still hung, kill the login process:

```
# /bin/kill -9 711
# /bin/ps aux | /bin/grep reid
root      4056    64 L    68 22 Slp     0:00  grep reid
```

If a process running a program refuses to be killed, you can at least recover the terminal by killing the login process. However, you cannot recover the process slot until you reboot the system.

## 5.2 FILE SYSTEM

The command /bin/df without arguments displays the current state of all file systems. The display includes the device name of the file system, the size of the file system in kilobytes, kilobytes used, kilobytes still available, the percentage of the file system used, and the mount point of the file system. Of particular interest is the percentage of the file system in use; if this figure is greater than 90%, you should consider taking steps to clean up the file system.

The following example demonstrates the command. Interpret (for example) the /usr line (last line) as follows: the device name of the file system is /dev/xdv01x (the x partition on controller 0, unit 1); the partition size is 326 megabytes; 277 megabytes are used; 16 megabytes are available; 94% of the file system is used; and the file system is mounted at /usr.

```
# /bin/df
Filesystem          kbytes    used   avail capacity  Mounted on
/dev/xdv00r          36155   30428    2111    94%    /
/dev/xdv01x         326291  277382   16279    94%    /usr
/dev/xdv20h         362459  323404    2809    99%    /usr2
/dev/xdv10h         362459  325724     489   100%    /general
/dev/xdv11s         150643  128266    7312    95%    /doc
/dev/xdv22h         362459  206764  119449    63%    /mfg
/dev/xdv21l         289331  208793   51604    80%    /usenet
/dev/xdv32h         355595  196487  123548    61%    /sources
```

The command /bin/du displays the disk usage of directories in kilobytes starting at a specified directory or the current working directory. This display helps you pinpoint the major users of disk space in a file system. You might perform a reverse numeric sort on the du output to better analyze the result:

```
# cd /doc
# /bin/du > /usr/adm/du_doc
# /usr/bin/sort -rn < /usr/adm/du_doc
128274  .
127686  ./usr
39462   ./usr/docmgr
29625   ./usr/bob
28765   ./usr/docmgr/desktop
26561   ./usr/burdette
15975   ./usr/docmgr/desktop/Concentrix.cab
11434   ./usr/burdette/lmath
10722   ./usr/burdette/desktop
9813    ./usr/chuck
8705    ./usr/bob/desktop
6559    ./usr/chuck/desktop
4778    ./usr/chuck/desktop/commands_V_5.drw
4430    ./usr/chuck/desktop/commands_V_5.drw/E.boo
4111    ./usr/docmgr/man
3848    ./usr/bob/kman
3708    ./usr/bob/desktop/admin.drw
```

The command /bin/ls -lgR displays information on individual files, including sizes and modification dates, starting at a specified directory or the current working directory. Note that file sizes are in bytes not kilobytes. This display helps you pinpoint individual files that are very large or very old.

Files that are likely candidates for deletion, archiving, or pruning (see Chapter 4) include the following:

- **System files that grow** – See the list of "System Files That Grow" in Chapter 1. Archive (as desired) and reduce the sizes of these files periodically.

- **Core dumps** – Programs that fail during execution may save the contents of memory in a file named core in the current working directory. These files are very large. In general, users who want to save core dumps rename them, and files named core are fair game for arbitrary deletion.

- **Temporary files** – If you encounter large files with names such as temp, tmp, foo, and bar, you might query the owners as to their utility. Such files are prime candidates for deletion. Files and subdirectories in the directory /tmp are candidates for deletion after a certain (usually short) period of time.

- **Object, executable, and listing files** – In development directories, these files are typically not necessary after linking and debugging. (They can be generated from the source files.) Arbitrarily deleting such files is not advised, but users should be counseled to clean up development directories.

The command /usr/bin/find provides an automated mechanism for applying a set of actions to selected files and is useful for cleaning up directories. The following script deletes files named core and a.out, files prefixed by a pound sign (#), and files suffixed by a tilde (~) in the current working directory and all subdirectories:

```
find "${@-.}" '(' \
        -name core -o -name a.out -o -name '#*' \
      -o -name '*~' \
 ')' -print -exec rm -f '{}' ';'
```

### 5.3.5.7 Implicit Migration to IP

A process may implicitly migrate to an IP to perform a system service. In particular, code performing an input/output operation must execute on the IOM to which the input/output device is attached. If the IP is not needed for another process, a process may continue to execute on the IP after the system service is complete.

You can prevent extraneous migration time on an IP by executing the process in a class that does not exist on the IP. For example, if you set the IP scheduling vectors for classes 3 and 4, and the CE scheduling vectors for 2 and 3:

```
setsched IP  0 -t  100 4 3
setsched IP  1 -t  100 4 3
setsched CL  0 -tc 100 1 3
setsched CE 0  -t  30  2 3
setsched CE 1  -t  30  2 3
setsched CE 2  -t  30  2 3
setsched CE 3  -t  30  2 3
setsched CE 4  -t  30  2 3
setsched CE 5  -t  30  2 3
setsched CE 6  -t  30  2 3
setsched CE 7  -t  30  2 3
setsched CE 8  -t  30  2 3
setsched CE 9  -t  30  2 3
setsched CE 10 -t  30  2 3
setsched CE 11 -t  30  2 3
setsched CE 12 -t  30  2 3
setsched CE 13 -t  30  2 3
setsched CE 14 -t  30  2 3
setsched CE 15 -t  30  2 3
```

And execute a process in class 2:

```
% /usr/alliant/execute -cl2 /usr/hacker/stay_on_ce
```

The process will immediately return to a CE after implicitly migrating to an IP to perform a system service.

You should not delete class 3 or class 4 from the scheduling vectors for the IPs. Such an action might cause the system to hang.

## 5.3.6 Real-time Scheduler

Concentrix-2800 has a real-time scheduler that provides a fast, deterministic means for scheduling work. The real-time scheduler requires a license for the FX/RT-2800 product and is explained in separate documentation.

## 5.4 CRONTAB

The command /etc/cron (clock daemon), which is started by the system start-up file /etc/rc, executes specified commands at specified times. The cron facility is a useful device for specifying periodic maintenance functions.

The commands to be executed by cron and the times of execution are specified in the plain text file /usr/lib/crontab. Each line of the file specifies one command in six fields separated by spaces or tabs:

1. Minute of execution (0–59).
2. Hour of execution (0–23).
3. Day (of the month) of execution (1–31).
4. Month of execution (1–12).
5. Day (of the week) of execution (1–7, where 1 means Monday).
6. Name of user under whose user identification and permissions the command is executed.
7. The command line to be executed. You can specify standard input to the command by following the command with a percent sign and text for each line of input.

In a time specification: an asterisk means all legal values, a list of comma-separated values means all values in the list, and two hyphen-separated values mean all values in the specified range.

The default /usr/lib/crontab file distributed with the system is as follows:

```
0 * * * *       root    (/bin/echo "^M";/bin/echo "    `/bin/hostname`
/bin/date` `/bin/echo "^M"` ") > /dev/console
0,10,20,30,40,50 * * * * root /etc/dmesg - >>/usr/adm/messages
0,15,30,45 * * * *      root /usr/lib/atrun
#00 1 * * *     daemon  /usr/new/lib/news/daily
#0 0 * * *      uucp    /usr/local/lib/uucp.daily
#7 5,12,18 * * *        uucp    /usr/local/lib/uucp.6hours
0 2 * * *       root    /bin/sh /usr/adm/daily 2>&1 | mail root
30 3 * * 6      root    /bin/sh /usr/adm/weekly 2>&1 | mail root
30 5 1 * *      root    /bin/sh /usr/adm/monthly 2>&1 | mail root
```

The fourth line, for example, executes /usr/lib/atrun (executes jobs for the at facility) every 15 minutes.

Do not remove the comments (initial pound signs) from the uucp lines unless you install UUCP.

You should examine and edit as desired the scripts executed in the last three lines.

## 5.5 KERNEL VARIABLES

You can patch the Concentrix-2800 kernel to modify operating system variables. See Patch_kernel in Appendix B for instructions and a list of variables. Follow the instructions carefully; do not bypass the steps for backing up the kernel and logging the changes. Resetting kernel variables arbitrarily can inhibit performance or make the system inoperative.

You can modify the Concentrix-2800 kernel at the source level. See Build_kernel in Appendix B.

# CHAPTER 6
# PRINT AND BATCH QUEUES

### Administration Files

| name | description |
|---|---|
| /dev/printer | socket for local print requests |
| /dev/* | print devices |
| /etc/printcap | printer capability data base ** (edit) |
| /etc/qconf | batch configuration file ** (edit) |
| /etc/ttys | terminal initialization data ** (edit) |
| /usr/bin/fpr | line printer Fortran filter |
| /usr/lib/lpf | line printer output filter |
| /usr/spool/batch/* | spooling area (queue) for batch jobs |
| /usr/spool/lpd/* | spooling area (queue) for print jobs |

** See *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

### Administration Commands (Appendix A)

| name | description |
|---|---|
| batchd | batch daemon |
| lpc | line printer control program |
| lpd | line printer daemon |
| netdae | batch network server daemon |
| netrecv | batch network receive daemon |

### Administration Procedures (Appendix B)

| name | description |
|---|---|
| Add_printer | add printer to system |
| Install_batch | install batch system |

**General Commands (Concentrix-2800 Commands and Applications)**

| name | description |
|---|---|
| lpq | spool queue examination program |
| lpr | print files on line printer |
| lprm | remove jobs from line printer queue |
| pr | format and display file |
| print | pr to the line printer |
| qdel | delete batch request |
| qsub | submit batch request |
| qstat | show batch request status |

## 6.1 PRINTERS AND PRINT QUEUES

The printer system permits the submission of print jobs to print queues (also called spooling areas) and the subsequent printing of the jobs on printers associated with the queues. Job submission is immediate – an interactive user can submit a print job and continue with other work. Job execution (the actual printing) is serial within a queue on a first-in first-out basis (unless the system administrator invervenes with the topq command of lpc).

The default system configuration consists of one printer named lp, an associated queue named /usr/spool/lpd, and an associated device named /dev/mlp0 – major number 0x90, minor number 0x1. In addition, /dev/lp is a soft link to /dev/mlp0. The device is a parallel printer connected to an Interphase controller.

You can attach additional parallel printers, raster output devices, and laser printers by adding appropriate controllers. You can attach serial printers to spare tty ports on terminal multiplexors. To queue jobs to additional printers, you must set up associated printer names, queues, and devices.

You can also queue print jobs to remote systems. Remote spooling via the network is handled with two print queues, one on the local system and one on the remote system. A remote job is queued locally and a daemon oversees its transfer to the remote system. If the remote system is unreachable, the job remains queued locally until the transfer is possible.

The following figure illustrates various printer capabilities.

**Sample Printer Configuration**

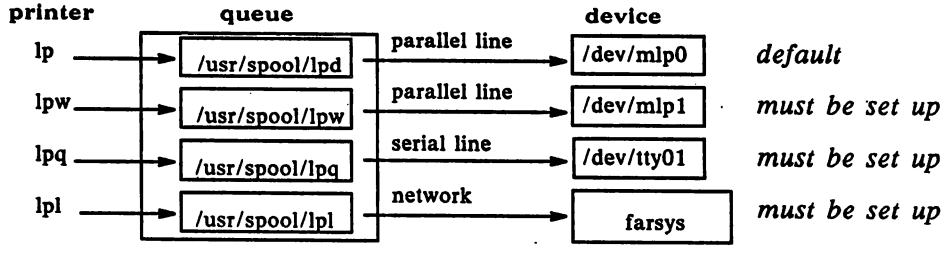

### 6.1.1 User Capabilities

A user submits and controls print jobs through the following commands:

- **lpr** – (print files on line printer) By default, this command submits a named file or files as a print job for the printer lp. For example:

  ```
  % lpr test1 test2
  ```

  You can specify an alternate printer by first setting the environment variable PRINTER:

  ```
  % setenv PRINTER lpw
  % lpr test1 test2
  ```

  Or by specify the -P option to lpr:

  ```
  % lpr -Plpw test1 test2
  ```

- **print** – (pr to the line printer) The -p option of lpr filters print jobs on input through the pr command, which divides the text into pages with headers and footers. The print command is equivalent to lpr with the -p option (or the pr command with output piped to lpr). The following commands are equivalent:

  ```
  % lpr -p test1 test2
  % print test1 test2
  % pr test1 test2 | lpr
  ```

- **lpq** – (spool queue examination program) This command lists the jobs that are in a print queue by priority, owner, job number (assigned by the system), file name, and size in bytes. Use the PRINTER environment variable or the -P option to specify other than the default queue.

- **lprm** – (remove jobs from line printer queue) This command removes a pending job (specified as the job number shown by lpq) or all jobs for the user. Use the PRINTER environment variable or the -P option to specify other than the default queue.

## 6.1.2 General Administration

Use the lpq command to examine queues. Note the job numbers and user names. The ranking indicates the order in which the jobs in the queue will be printed. The size (in bytes) gives you an idea of the relative sizes of the jobs in the queue; a rule of thumb is one page per 4096 bytes for plain text. Note the following example:

```
# lpq
Rank    Owner          Job     Files                          Total Size
1st     hacker         747     h1234.f                        2040 bytes
2nd     windy          748     memo                           580124 bytes
3rd     windy          749     memo2                          623100 bytes
4th     hacker         750     h1235.f                        5004 bytes
```

As superuser, you can remove any user's jobs with lprm. Specify lprm *job-number* to remove discrete jobs. The following command removes jobs 748 and 749 from the default queue:

```
# lprm 748 749
```

Specify lprm *user-name* to remove all jobs belonging to specified users. The following command removes all jobs in the default queue belonging to user windy:

```
# lprm windy
```

Specify lprm – to remove all jobs in a queue:

```
# lprm -
```

### 6.1.2.1 Line Printer Control (lpc) Program

The administration command /etc/lpc provides status information on and permits additional control over printers and print queues. The following command prints the status of lp:

```
# /etc/lpc status lp
lp:
        queuing is enabled
        printing is enabled
        3 entries in spool area
        waiting for lp to become ready (offline ?)
```

If no jobs are active or pending, the status should be as follows:

```
# /etc/lpc status lp
lp:
        queuing is enabled
        printing is enabled
        no entries
        no daemon present
```

You can run /etc/lpc in command mode (as above) or interactive mode. You enter interactive mode by not specifying any arguments; use the lpc help command to display brief descriptions of the lpc commands; use the lpc quit command to return to shell level. Note the following example:

```
# /etc/lpc
lpc> help
Commands may be abbreviated. Commands are:

abort   enable  disable quit    start   stop
clean   exit    help    restart status  topq
lpc> help topq
topq    put job at top of printer queue
lpc> status
lp:
        queuing is enabled
        printing is enabled
        no entries
        no daemon present
lpc> quit
```

The /etc/lpc commands provide the following capabilities:

- **Job submission** – To prevent users from submitting additional jobs:

  ```
  disable all
  disable printer ...
  ```

  To enable users to once again submit jobs:

  ```
  enable all
  enable printer ...
  ```

- **Job printing** – To stop printing jobs after the current job completes:

  ```
  stop all
  stop printer ...
  ```

  To abort the current print job and stop printing jobs:

  ```
  abort all
  abort printer ...
  ```

  To start printing jobs again:

  ```
  start all
  start printer ...
  ```

- **Restart** – To start printing jobs again if the print daemon dies prematurely:

  ```
  restart all
  restart printer ...
  ```

- **Remove jobs** – To remove all jobs (removes all files in the spooling directory starting with cf, tf, and df):

  ```
  clean all
  clean printer ...
  ```

- **Change rank** – To change the printing order of jobs:

  topq *printer* [ *job-number* ... ] [ *user-name* ... ]

  The specified jobs or the jobs belonging to the specified users are placed at the top of the queue in the order indicated.

### 6.1.2.2 Trouble-shooting and Error Messages

If the printing device refuses to print and the associated queue contains jobs, take the following steps:

1. Check the hardware. Ensure that the printing device is online; users often forget to reset the device after stopping it to remove listings. Ensure that the printing device has paper and does not have an error condition.

2. Check the queue. If the hardware is correct, execute lpq on the printer. If the queue contains jobs but no print daemon is present, restart the queue with the restart command of lpc. Also examine the file specified by lf in /etc/printcap, usually /usr/adm/lp*printer*-errs, to see if any error messages are logged.

3. Kill an offensive job. If the queue contains jobs and the print daemon is active, or the previous step fails, the problem is typically a hang on the current job. A hang usually is the result of a user attempting to print a binary file, although an ASCII file can occasionally cause a hang. Remove the job with lprm, or abort the job and start the queue again with the abort and start commands of lpc.

4. If the problem persists, check the filter specifications for the printer. A bad filter specification may cause a job requesting that filter to hang.

Expected hardware conditions that stop the printer are running out of paper and using up the print ribbon. Both conditions are remedied in a short time by following the manufacturer's instructions. Other conditions may require a service call.

If the hardware is expected to be out of service for some time, you should disable the associated queue with the disable command of lpc. When the hardware comes back on line, enable the queue with the enable command of lpc.

### 6.1.2.3 Error Messsages (lpc)

`could't start printer`

See "cannot start daemon" under the lpr errors below.

`cannot examine spool directory`

Error messages beginning with "cannot ..." are usually due to incorrect ownership and/or protection mode of the lock file, spooling directory, and/or the lpc program.

### 6.1.2.4 Error Messsages (lpq)

`waiting for` *printer* `to become ready (offline ?)`

Typically, the printer is offline or out of paper, or a hardware problem occurred. Another possible cause is that another process, such as an output filter, has an exclusive open on the device; your only recourse is to kill the offending program(s) and restart the printer with lpc.

*printer* `is ready and waiting`

If the daemon is hung (jobs are in the queue but do not print), use lpc to abort the current daemon and start a new one.

`waiting for` *host* `to come up`

A daemon is trying to connect to the remote system named *host* to send files from the local queue. If the remote system is up, lpd on the remote machine is probably dead or hung and should be restarted.

`sending to` *host*

The files should be in the process of being transferred to the remote *host*. If not, the local daemon should be aborted and started with lpc.

`Warning:` *printer* `is down`

The printer is unavailable.

`Warning: no daemon present`

The lpd process overseeing the spooling queue does not exist. Normally occurs only when the daemon has unexpectedly died. The error log file for the printer should be checked for a diagnostic from the deceased process. Use the restart command of lpc to restart the daemon.

### 6.1.2.5 Error Messsages (lpr)

`lpr: printer: unknown printer`

Indicates a typing error or an incorrect entry in /etc/printcap.

`lpr: printer: jobs queued, but cannot start daemon`

Usually means that the printer server started at boot time has died or is hung. Check for the existence of the local socket /dev/printer (if it does not exist, no lpd process is running). Use the following command to list the process identifiers of the running lpd processes.

```
# ps ax | fgrep lpd
```

Kill the master daemon, if it exists, using the following command. The master daemon is the lpd process whose process identification is not listed in any lock file in a spooling directory under /usr/spool.

```
# kill -9 pid
```

Remove /dev/printer (if it exists) and restart the daemon (and printer):

```
# rm /dev/printer
# /usr/lib/lpd
```

Another possible error is that the lpr program does not have set owner and set group permissions with ownership of root and group of daemon. The permissions and ownership should appear as follows. Use /etc/chown, chgrp, and chmod to adjust ownership and permissions as necessary.

```
# ls -lg /usr/ucb/lpr
-rws--s--x  1 root     daemon         49152 Oct 30   1987    /usr/ucb/lpr
```

`lpr: printer: printer queue is disabled`

The queue was turned off with the disable command of lpc. Use the enable command of lpc to permit jobs to be queued again.

### 6.1.2.6 Error Messsages (lprm)

`lprm: printer: cannot restart printer daemon`

See "cannot start daemon" under the lpr errors above.

### 6.1.3 Default Printer

No actions are required to set up the default line printer (other than physically connecting the printer to the controller). The default line printer is defined as the only entry in the /etc/printcap file that is distributed with the system.

The /etc/printcap file defines the printer capabilities for each printer on the system. The file as distributed has the following entry:

```
lp|local line printer:\
        :lp=/dev/lp:sd=/usr/spool/lpd:lf=/usr/adm/lpd-errs:fo:\
        :rf=/usr/bin/fpr:of=/usr/lib/lpf:pw#80:
```

The first field of the entry contains the name of the printer (lp) and a description (local line printer). The remaining fields define the printer as follows:

- lp – The print device is defined as /dev/lp. In the system as distributed, /dev/lp is a soft link to /dev/mlp0 – major number 0x90, minor number 0x1. The files are owned by root, with group wheel, and permissions 600.
- sd – The spooling directory (print queue) is defined as /usr/spool/lpd. This directory exists in the system as distributed owned by daemon, with group daemon, and permissions 775.
- lf – The file for error logging is defined as /usr/adm/lpd-errs.
- fo – The print device is defined to print an initial form feed.
- rf – The Fortran filter is defined as /usr/bin/fpr.
- of – The output filter is defined as /usr/lib/lpf.
- pw – The page width is defined as 80 characters.

### 6.1.4 Additional Parallel Printers

To add another parallel printer, first have the appropriate controller installed in the machine and physically connect the printer. Then create a special file for the device. The major number is 0x90. The minor number is 0x*cc*0001 where *cc* is the controller number.

Permissions should be 600 to prevent users from writing to the device. The following example creates a special file for a second printer:

```
# /etc/mknod /dev/mlp1 c 0x90 0x10001
# chmod 600 /dev/mlp1
```

Make a spooling directory (print queue) under /usr/spool. The ownership should be daemon, the group daemon, and the permissions 775. The following example makes a spooling directory named lpw:

```
# cd /usr/spool
# mkdir lpw
# chown daemon lpw
# chgrp daemon lpw
# chmod 775 lpw
```

Edit /etc/printcap to add an entry for the new printer. The simplest technique is to copy the entry for lp, paste it, and modify it. The following example demonstrates an /etc/printcap entry for a printer named lpw:

```
lpw|local wide-paper line printer:\
        :lp=/dev/mlp1:sd=/usr/spool/lpw:lf=/usr/adm/lpw-errs:fo:\
        :rf=/usr/bin/fpr:of=/usr/lib/lpf:pw#132:
```

The entries for af (if used), lf, lp, and sd should be different for each entry in /etc/printcap. In the example, a printer named lpw queues its jobs in /usr/spool/lpw and prints to the device /dev/mlp1. The page width is 132.

### 6.1.5 Serial Printers

To add a serial printer, first physically connect the printer to a tty port with a standard terminal-multiplexor cable; see "Terminal Devices" in Chapter 3. Ensure that a special file exists for the port and that the file is set up as direct connect (not a modem). The major number is 0x98. The minor number is 0x*c*000*u*, where *c* is the controller number (0x0 through 0x7) and *u* is the port number (0x0 through 0xf). The name of the device is typically /dev/tty*cu*, where *c* is the controller number and *u* is the port number in hexadecimal.

Permissions should be 600 to prevent users from writing to the device. The following example sets up tty port 15 of controller 0 for a serial printer:

```
# rm /dev/tty0f
# /etc/mknod /dev/tty0f c 0x98 0xf
# chmod 600 /dev/tty0f
```

Create an entry (or adjust the existing entry) in /dev/ttys for the device. The on flag must be absent or specified as off. The argument for terminal characteristics is usually 6, 7, or 2 (direct connect 2400-baud, 4800-baud, or 9600-baud line). The baud setting must match the setting for the device. The following example shows an entry for /dev/tty0f:

```
tty0f   "/etc/getty std.9600" vt100   off
```

Make a spooling directory (print queue) under /usr/spool. The ownership should be daemon, the group daemon, and the permissions 775. The following example makes a spooling directory named lpq:

```
# cd /usr/spool
# mkdir lpq
# chown daemon lpq
# chgrp daemon lpq
# chmod 775 lpq
```

Edit /etc/printcap to add an entry for the new printer. The simplest technique is to copy the entry for lp, paste it, and modify it. The following example demonstrates an /etc/printcap entry for a printer named lpq:

```
lpq|local letter quality serial printer:\
        :lp=/dev/tty0f:sd=/usr/spool/lpq:lf=/usr/adm/lpq-errs:fo:\
        :br#9600:fs#06320:\
        :rf=/usr/bin/fpr:of=/usr/lib/lpf:pw#80:
```

The entries for af (if used), lf, lp, and sd should be different for each entry in /etc/printcap. For a serial printer, you must also specify the br entry (baud rate), and the fs and/or fc entry (set/clear sg flags; see the *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: tty). In the example, a printer named lpq queues its jobs in /usr/spool/lpq and prints to the device /dev/tty0f at a baud rate of 9600 with a flags setting of 06320 (CRMOD on, no parity, and XTABS on).

### 6.1.6 Remote Printers

To enable printing to a printer on another system, first make a spooling directory (print queue) under /usr/spool. The ownership should be daemon, the group daemon, and the permissions 775. The following example makes a spooling directory named lpl:

```
# cd /usr/spool
# mkdir lpl
# chown daemon lpl
# chgrp daemon lpl
# chmod 775 lpl
```

Edit /etc/printcap to add an entry for the new printer. The simplest technique is to copy the entry for lp, paste it, and modify it. The following example demonstrates an /etc/printcap entry for a printer named lpl:

```
lpl|remote laser printer:\
        :lp:rm=farsys:rp=lpl:sd=/usr/spool/lpl:
```

For a remote printer, the lp entry must be null, an rm entry must specify the name of the remote system, and an rp entry must specify the name of the printer on the remote system. Typically, only these entries and the sd entry are specified. In the example, a printer named lpl queues its jobs in /usr/spool/lpl and prints to the printer named lpl on the system named farsys.

## 6.1.7 Output Filters

You can develop output filters or use one of several that are supplied with the system to handle device dependencies and to perform accounting functions. Various fields in /etc/printcap specify the names of the output filters, as shown in column 1 of the following table. The of and if filters, when specified in /etc/printcap, apply to every job queued to the printer. The remaining filters require that the user specify an argument to lpr as shown in the second column. If the of filter is specified alone, the of filter processes an entire job. If the of filter and another filter are specified, the of filter is called at the beginning of the job and the other filter is called for each file in the job.

**Output Filters**

| printcap | user | description |
|---|---|---|
| of= | automatic | Standard output filter. |
| if= | automatic | Accounting filter. |
| cf= | lpr -c | Cifplot filter. |
| df= | lpr -d | Tex filter. |
| gf= | lpr -g | Plot filter |
| nf= | lpr -n | Ditroff filter. |
| rf= | lpr -f | Fortran filter. |
| tf= | lpr -t | Troff filter. |
| vf= | lpr -v | Raster filter. |

All filters except the output filter are capable of writing to the accounting file. The accounting file is specified by the af entry in /etc/printcap.

The system as distributed contains a generic output filter named /usr/lib/lpf, a Fortran filter named /usr/bin/fpr, and various other filters in /usr/lib. The lp entry in /etc/printcap as distributed specifies of=/usr/lib/lpf and rf=/usr/bin/fpr. The default configuration, then, invokes /usr/lib/lpf for each job, and /usr/bin/fpr for each file if lpr -f is specified.

The line printer daemon (/usr/lib/lpd) spawns filters with standard input attached to the data to be printed, and standard output attached to the printer. Standard error is attached to the file specified by lf. The line printer daemon calls a filter using one of three sets of arguments depending on the filter type (output filter, accounting filter, any other filter):

*of-filter* **-w***width* **-l***length*

*if-filter* [**-c**] **-w***width* **-l***length* **-i***indent* **-n** *login* **-h** *host accounting-file*

*other-filter* **-w***width* **-l***length* **-n** *login* **-h** *host accounting-file*

The *width* and *length* values are the pw and pl entries in /etc/printcap. The -c argument is only supplied when control characters are to be passed uninterpreted to the printer (when the -l option of lpr is used to print the file). The -n and -h arguments specify the login name and host name of the job owner. The last argument is the name of the accounting file, the af entry in /etc/printcap.

A filter must return one of the following exit codes:

- 0 – No errors.
- 1 – Job should be reprinted.
- 2 – Job should be thrown away.

When lprm sends a kill signal to the lpd process controlling printing, a SIGINT signal is propagated to all filters and descendents of filters. This signal can be trapped by filters that need to perform cleanup operations such as deleting temporary files.

### 6.1.8 Printer Accounting

Printer accounting is not on by default. To turn printer accounting on, you must change the entry for the printer in /etc/printcap: specify the if option instead of the of option, using the same filter name; specify the af option to identify the accounting file (typically named after the printer in /usr/adm). The following example adjusts /etc/printcap to turn on accounting for the standard job filter:

```
lp|local line printer:\
        :lp=/dev/lp:sd=/usr/spool/lpd:lf=/usr/adm/lpd-errs:fo:\
        :rf=/usr/bin/fpr:if=/usr/lib/lpf:af=/usr/adm/lpd-acct:pw#80:
```

Touch the accounting file to make sure it exists, as shown in the following example:

```
# touch /usr/adm/lpd-acct
```

Kill the line printer daemon and restart it to force it to reread the /etc/printcap data base, as shown in the following example:

```
# /etc/lpc abort lp
# /etc/lpc start lp
```

## 6.2 BATCH QUEUES

The batch system permits the execution of command scripts as jobs on a priority (default) or round-robin basis as resources become available. You can administer the batch system to regulate the number and types of command scripts that can execute concurrently. Command scripts can also be transferred for execution on other machines.

A batch job (that is, the command script to be executed) is submitted to a queue. Associated with each queue are one or more logical devices; a device can be shared by more than one queue. Batch jobs are processed from a queue in order of priority (default) or on a round-robin basis using as many logical devices as are available or necessary.

The following diagram demonstrates a system with four queues. The first queue, named batch, is the default queue for job submission and can process two jobs concurrently; additional jobs wait. The queue batch@farsys transmits submitted jobs to the queue named batch on the system farsys. The queue blittle processes up to three jobs at a time and the queue bbig processes no more than one job at a time. The allocation of logical devices on the example system permits a maximum of four jobs to execute concurrently on the local system: two shared between batch and blittle; one dedicated to blittle; and one dedicated to bbig.

**Sample Batch Configuration**

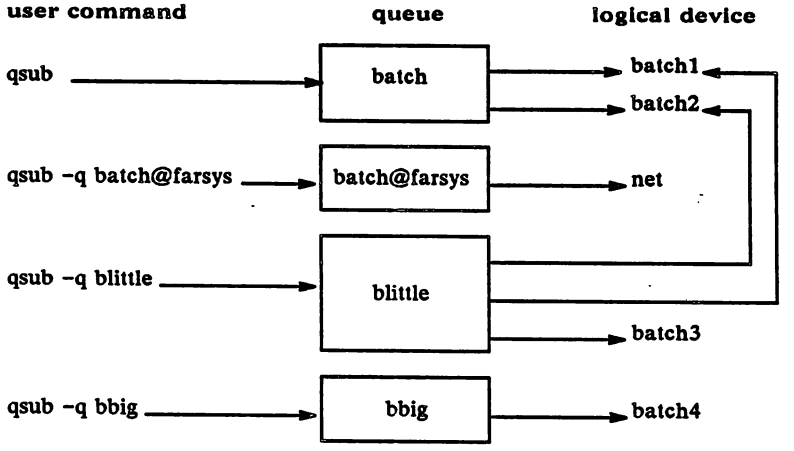


The Concentrix-2800 batch system is based upon the Multi-Device Queueing System (MDQS) developed at the U.S. Army's Ballistic Research Laboratory by D. P. Kingston. For a complete description of MDQS refer to: Kingston, Douglas P. and Muuss, Michael J., *The Multi-Device Queueing System*, 1982 Summer USENIX Conference.

### 6.2.1 Batch Setup and Queue Configuration

The system as distributed contains the file /etc/qconf. This file must be modified for the batch system to operate. The site-specific start-up file should be correct as distributed but should be checked to ensure that the portion pertaining to batchd and netdae is not commented out. The procedure Install_batch in Appendix B provides detailed steps.

You *must* make the following modifications to /etc/qconf:

- **hostname** – The hostname value on line 4 must be changed from localname to the name of your system.
- **batch@remote** – The two lines containing batch@remote and the net line must be commented out. Or the word remote (three occurrences) must be changed to the name of another system with batch queues (and the word batch must be changed to an actual queue name if the system is not running a queue named batch).

The modifications shown above enable two local devices named batch1 and batch2 mapped to a local queue named batch. And optionally a network device named net mapped to a remote queue.

The /etc/qconf file contains three sections: Batch Device Descriptions; Batch Queue Descriptions; and Batch Mappings. These sections can be modified to add, delete, and change the names of logical devices and queues, and to change the mappings of devices to queues. See the procedure Install_batch in Appendix B.

Do not modify the last line of the file.

The start-up script /etc/rc.local contains lines that clean out the batch spooling area, start the batch daemon, and start batch network server daemon. If you start the daemons explicitly, invoke them as background jobs:

```
# /usr/lib/batch/batchd &
# /usr/lib/batch/netdae -n 2 /usr/lib/batchd/netrecv &
```

### 6.2.2 User Capabilities and General Administration

The mechanism for submitting jobs is the /usr/alliant/qsub command. The -q option specifies the queue (defaults to batch). The -p option establishes the priority (0–127, low to high; defaults to 64). The -a option delays execution until after a specified time. The following example submits a command script named job to the queue named batch:

```
% qsub job
Request Q00116.10 (hacker.10) Queued
```

The next example submits a job named bigjob to the queue named bbig with a low priority:

```
% qsub -q bbig -p 1 bigjob
Request Q00116.11 (hacker.11) Queued
```

The job is identified by the name of the submitting user suffixed by a decimal point and an integer.

For local jobs, the user's environment and current working directory are saved with the request and restored when the job runs. For remote jobs, the working directory defaults to /tmp when the job runs on the remote system. Therefore, a script destined for remote execution should change directory (cd command) in the shell script.

The /usr/alliant/qstat command displays the status of the batch system in three parts: active device – jobs currently being executed, by logical device; request queue – jobs ready but waiting to be executed, by queue; delay queue – jobs on hold to a specified time, by queue. The -al options display all jobs in long format.

```
# /usr/alliant/qstat -al

ACTIVE-DEVICE  REQUEST       FILE        PRI    SIZE    SUBMITTED
batch1         hacker.20     bj2         64     10044   Jul 25 17:46
batch2         hacker.19     bj1         64     1224    Jul 25 17:46
batch3         [empty]
batch4         [empty]
net            [empty]

REQUEST-QUEUE  REQUEST       FILE        PRI    SIZE    SUBMITTED
batch          hacker.21     bj3         64     14308   Jul 25 17:57
batch          hacker.22     bj4         64     12244   Jul 25 17:58
blittle        [empty]
bbig           [empty]
batch@farsys   [empty]

DELAY-QUEUE    REQUEST       FILE        PRI    SIZE    SUBMIT-TIME
batch          hacker.17     bj0         64     1564    Jul 25 18:00
batch          hacker.18     bj00        64     1932    Jul 25 18:00
```

The qstat command does not work for jobs on remote systems. However, you can place it in an rsh command. For example:

```
# /usr/ucb/rsh farsys /usr/alliant/qstat -al
```

The /usr/alliant/qdel command deletes a batch request, which must be identified by the request name shown in column 2 of the long qstat display. Use the -k option to delete active jobs; otherwise, omit the -k option. The following example deletes two jobs still in a queue:

```
# /usr/alliant/qdel hacker.17 hacker.18
```

The next example deletes one active job:

```
# /usr/alliant/qdel -k hacker.19
```

You must be superuser to delete a job other than your own.

The qdel command does not work for jobs on remote systems. However, you can place it in an rsh command. For example:

```
# /usr/ucb/rsh farsys /usr/alliant/qdel -k hacker.19
```

You cannot direcly modify the priority or other characteristics of a job. The recommended procedure is to delete the job and resubmit it.

The *Concentrix-2800 Commands and Applications Manual* documents the qsub, qstat, and qdel commands.

# CHAPTER 7
# NETWORKS

**Administration Files**

| name | description |
|---|---|
| /etc/ethers | Ethernet address data base (edit) |
| /etc/exports | NFS file systems being exported ** (edit) |
| /etc/ftpusers | list of users excluded by ftpd (edit) |
| /etc/gateways | routing daemon data base ** (edit) |
| /etc/gettytab | terminal configuration data base ** (edit) |
| /etc/group | group file (edit) |
| /etc/hosts | host name data base ** (edit) |
| /etc/hosts.equiv | list of trusted hosts ** (edit) |
| /etc/inetd.conf | inetd configuration file (edit) |
| /etc/networks | network name data base ** (edit) |
| /etc/passwd | password file ** (edit) |
| /etc/protocols | protocol name data base ** (edit) |
| /etc/publickey | public key database ** |
| /etc/rc.local | site-specific command script for starting Concentrix (edit) |
| /etc/remote | remote host description file ** (edit) |
| /etc/rmtab | remotely mounted file system table ** |
| /etc/rpc | rpc program number data base ** (edit) |
| /etc/servers | inet server data base ** (edit) |
| /etc/services | service name data base ** (edit) |
| /etc/ttys | terminal initialization data ** (edit) |
| /etc/yp/updaters | makefile for updating NIS databases ** |
| /etc/yp/ypfiles | network information service files ** (edit) |
| ~/.netrc | user information for ftp ** (edit) |
| .rhosts | personal list of trusted hosts ** (edit) |
| /usr/lib/uucp/L.aliases | UUCP hostname alias file ** (edit) |
| /usr/lib/uucp/L.cmds | UUCP remote command permissions file ** (edit) |
| /usr/lib/uucp/L-devices | UUCP device description file ** (edit) |
| /usr/lib/uucp/L-dialcodes | UUCP phone number index file ** (edit) |
| /usr/lib/uucp/L.sys | UUCP remote host description file ** (edit) |
| /usr/lib/uucp/USERFILE | UUCP pathname permissions file ** (edit) |

** See *Concentrix-2800 System Reference:Miscellaneous Routines, Special Files, and File Formats*
(edit) can be edited by system administrator

## Administration Commands (Appendix A)

| name | description |
|---|---|
| automount | automatically mount NFS file system |
| comsat | biff server |
| ftpd | DARPA Internet file transfer protocol server |
| ifconfig | configure network interface parameters |
| ifstats | print out Ethernet board statistics |
| keyserv | daemon that stores private encryption keys for secure NFS |
| mailstats | print out sendmail statistics |
| makedbm | make a network information service dbm file |
| mconnect | connnect to SMTP mail server socket |
| mknod | build special file |
| mount | mount file systems |
| mountd | NFS mount request server |
| newkey | establish public keys for users/superusers on secure NFS/NIS |
| nfsd | NFS daemon |
| ntalkd | server for talk program |
| ping | send icmp echo packet (network debugging) |
| portmap | DARPA port to RPC program number mapper |
| rdate | set system date from a remote host |
| rdump | file system dump across the network |
| rexecd | remote execution server |
| rlogind | remote login server |
| route | manually manipulate the routing tables |
| routed | network routing daemon |
| rmt | remote magtape protocol module |
| rrestore | restore a file system dump across the network |
| rshd | remote shell server |
| rwalld | network rwall server |
| rwhod | system status server |
| sendmail | send mail over the internet |
| showmount | show all remote mounts |
| syslog | log system messages |
| telnetd | DARPA TELNET protocol server |
| tftpd | DARPA trivial file transfer protocol server |
| uucico | transfer files queued by uucp or uux |
| uuclean | UUCP spool directory clean-up |
| uucpd | UUCP daemon |
| uupoll | poll a remote UUCP site |
| uusnap | show snapshot of the UUCP system |
| uuxqt | UUCP execution file interpreter |

**Administration Commands (Appendix A), continued**

| name | description |
|---|---|
| ypbind | network information service binder process |
| ypinit | build and install network information service data base |
| ypmake | rebuild network information service data base |
| yppasswd | server for modifying network information service password file |
| yppoll | network information service get map version |
| yppull | network information service get map from master |
| yppush | network information service send map to slaves |
| ypserv | network information service server process |
| ypupdated | daemon that updates information in the NIS |
| ypxfr | transfer an NIS map to local host |
| ypwhich | determine networl information service master |

**Administration Procedures (Appendix B)**

| name | description |
|---|---|
| Add_ethernet | add system to Ethernet |
| Add_uucp | add system to UUCP network |
| Add_yp | add network information service to network |

**General Commands (Concentrix-2800 Commands and Applications Manual)**

| name | description |
|---|---|
| biff | be notified if mail arrives |
| chmod | change file access mode |
| domainname | set or display name of current domain system |
| ftp | file transfer program |
| hostname | set/display host system name |
| kill | terminate job or process |
| netstat | show network status |
| rcp | remote file copy |
| rlogin | remote login |
| rmail | handle remote mail received via uucp |
| rsh | remote shell |
| ruptime | show host status of local machines |
| rwall | write to all users over a network |
| rwho | show who's logged in on local machines |
| talk | talk to another user |
| telnet | user interface to the TELNET protocol |
| tftp | trivial file transfer program |
| tip | connect to a remote system |

**General Commands, continued**

| name | description |
|---|---|
| uucp | unix to unix copy |
| uulog | display UUCP log files |
| uuname | list names of UUCP hosts |
| uuq | examine or manipulate the UUCP queue |
| uusend | send file to remote host |
| uux | unix to unix command execution |
| ypcat | display values in a network information service data base |

**Special File Entries**

| | |
|---|---|
| arp | Address Resolution Protocol ** |
| icmp | Internet Control Message Protocol ** |
| if | general properties of network interfaces ** |
| inet | Internet protocol family ** |
| ip | Internet Protocol ** |
| pty | pseudo terminal driver ** |
| tty | general terminal interface ** |
| tcp | Internet Transmission Control Protocol ** |
| udp | Internet User Datagram Protocol ** |

** See *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*

## 7.1 LOCAL AREA NETWORKS

Concentrix-2800 supports local area networks through Ethernet hardware and Internet protocols. The hardware for one network consists of an Ethernet controller in each participating machine, a transceiver for each machine, drop cables connecting controllers to transceivers, and coaxial cables connecting transceivers. The following diagram illustrates a network of three computers. Multiplexors can be used instead of individual transceivers; multiplexors connect up to eight drop cables.

**Ethernet Local Area Network**

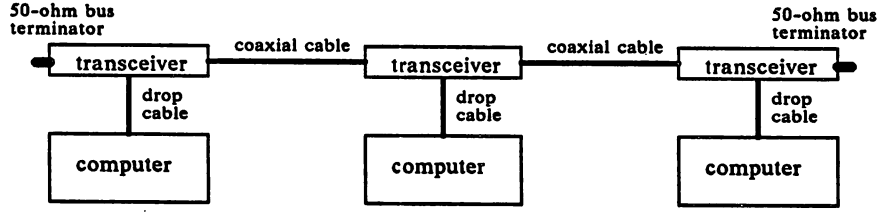


**The coaxial cable between transceivers must be at least 10 feet long. The drop cables must be no more than 10 feet long.**

The system as distributed supports the following Internet protocols:

| | |
|---|---|
| IP | Internet protocol |
| ICMP | Internet control message protocol |
| GGP | gateway-gateway protocol |
| TCP | transmission control protocol |
| PUP | PARC universal packet protocol |
| UDP | user datagram protocol |

Logically, a local area network appears as a number of host systems that can communicate with each other and with host systems on other networks through services such as login (remote login) and ftp (file transfer program). A network is identified by a name and an Internet address; each host system is identified by a host name and an Internet address that is an extension of the network Internet address. Ordinary users access systems through their host names and do not need to know the network names and Internet addresses.

The following figure illustrates two networks named n1net and n2net whose Internet addresses are 192.0.1 and 192.0.2, respectively. The unshaded boxes represent host systems; the first line in each box is the system name and the second line is the system's Internet number.

**Sample Networks**

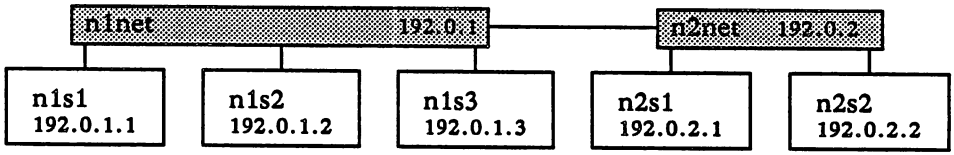


## 7.1.1 Internet Services

The system as distributed supports the following Internet services. Most of the servers (the programs implementing the services) are started as needed by the /etc/inetd daemon. The /etc/inetd daemon is started at system start-up by the script /etc/rc (as distributed). A few servers are started directly at system start-up.

- **biff** – Be notified if mail arrives, as with the biff command. Serviced by the /etc/comsat daemon, which is started as needed by the /etc/inetd daemon.
- **exec** – Remote execution, as with the rexec system routine. Serviced by the /etc/rexecd daemon, which is started as needed by the /etc/inetd daemon.
- **finger** – User information lookup program. Serviced by the /etc/fingerd daemon, which is started as needed by the /etc/inetd daemon.
- **ftp** – File transfer program. Serviced by the /etc/ftpd daemon, which is started as needed by the /etc/inetd daemon.
- **login** – Remote login, as with the rlogin command. Serviced by the /etc/rlogind daemon, which is started as needed by the /etc/inetd daemon.
- **mount** – Mount remote file systems, as with the mount –t nfs command. Serviced by the /usr/etc/rpc.mountd daemon, which is started as needed by the /etc/inetd daemon.
- **named** – Internet domain name server. Serviced by the /etc/named daemon, which is included but commented out of the system start-up script /etc/rc.local (as distributed). Uncomment the invocation if you want the daemon to start at system start-up.
- **ntalk** – Talk to another user, as with the talk command. Serviced by the /etc/ntalkd daemon, which is started as needed by the /etc/inetd daemon.
- **portmap** – DARPA port to RPC program number mapper. Serviced by the /etc/portmap daemon, which is started automatically at system start-up by the script /etc/rc.local (as distributed).

- **rexd** – Remote program execution. Serviced by the /usr/etc/rpc.rexd daemon, which is started as needed by the /etc/inetd daemon.

- **route** – Network routing, as with the /etc/route command. Serviced by the /etc/routed daemon, which is included but commented out of the system start-up script /etc/rc.local (as distributed). Uncomment the invocation if you want the daemon to start at system start-up.

- **rstatd** – Kernel statistics server. Serviced by the /usr/etc/rpc.rstatd daemon, which is started as needed by the /etc/inetd daemon.

- **rusersd** – Network username server. Serviced by the /usr/etc/rpc.rusersd daemon, which is started as needed by the /etc/inetd daemon.

- **shell** – Remote shell, as with the rsh command. Serviced by the /etc/rshd daemon, which is started as needed by the /etc/inetd daemon.

- **smtp** – Send mail over the internet. Serviced by the /usr/lib/sendmail –bd daemon, which is started automatically at system start-up by the script /etc/rc.local (as distributed).

- **sprayd** – Spray server. Serviced by the /usr/etc/rpc.sprayd daemon, which is started as needed by the /etc/inetd daemon.

- **syslog** – Log system messages. Serviced by the /etc/syslog daemon, which is started automatically at system start-up by the script /etc/rc.local (as distributed).

- **telnet** – User interface to the TELNET protocol, as with the telnet command. Serviced by the /etc/telnetd daemon, which is started as needed by the /etc/inetd daemon.

- **tftp** – Trivial file transfer program. Serviced by the /usr/etc/in.tftpd daemon, which is started as needed by the /etc/inetd daemon.

- **timed** – Time server. Serviced by the /etc/timed daemon, which is included but commented out of the system start-up script /etc/rc.local (as distributed). Uncomment the invocation if you want the daemon to start at system start-up.

- **wall** – Write to all users over a network, as with the rwall command. Serviced by the /usr/etc/rpc.rwalld daemon, which is started as needed by the /etc/inetd daemon.

- **who** – System status, as with the rwho and ruptime commands. Serviced by the /etc/rwhod daemon, which is started at system start-up by the script /etc/rc (as distributed).

The following networking commands are also supported:

- **netstat** – Show network status.
- **ping** – Send icmp echo packet (network debugging).
- **rdump** – File system dump across the network.
- **rmt** – Remote magtape protocol module.
- **rpcinfo** – Report RPC information.
- **rrestore** – Restore a file system dump across the network.

## 7.1.2 Device Names and Addresses

Certain administrative procedures require Ethernet controller names and Ethernet addresses. In addition, the login (remote login) service requires pseudo terminal devices.

### 7.1.2.1 Ethernet Controller Names

Ethernet controllers are not accessed through special files in /dev but do have system generated names. The name of an Ethernet controller may be ex*n* (Excelan), or en*n* (CMC) where *n* is a digit that depends on the controller's configuration in the system. On a system with one Ethernet controller and no other network controllers, the name of the Ethernet controller is typically either ex0 or en0.

Use the /usr/alliant/syscfg -d command to determine the name of an Ethernet controller (note the last line of the following example):

```
% syscfg -d
...

IP-0:
  xdv0 (port=0) (major= 96) (addr=0xfd00ee40) Xylogics 752/754 Disk Controller
         Drive(s) Online: 0 1
  xdv1 (port=0) (major= 96) (addr=0xfd00ee50) Xylogics 752/754 Disk Controller
         Drive(s) Online: 0 1
  xdv2 (port=0) (major= 96) (addr=0xfd00ee60) Xylogics 752/754 Disk Controller
         Drive(s) Online: 0 1
  xdv3 (port=0) (major= 96) (addr=0xfd00ee70) Xylogics 752/754 Disk Controller
         Drive(s) Online: 0 1
  xdv4 (port=0) (major= 96) (addr=0xfd00ee80) Xylogics 752/754 Disk Controller
         Drive(s) Online: 0 1
 scsi0 (port=0) (major=165) (addr=0xfd00c000)
  imb0 (port=0) (major=224) (addr=0x0)
   en0 (port=0) (major=206) (addr=0xfe000000)
```

### 7.1.2.2 Ethernet Addresses

Each connected system has a 6-byte Ethernet address. The address appears at system start-up near the end of the boot script as six bytes in hexadecimal format separated by periods. Consult the hard copy from your console printer or the file /usr/adm/messages (or reboot the system) to determine the Ethernet address. A sample from a boot script follows:

```
Jun 18 06:56:10 buddy vmunix: ip00> Paging File space: 38016 pages, 148.5 Mb
Jun 18 06:56:10 buddy vmunix: ip00> Swapping File space: 31103 pages, 121.5 Mb
Jun 18 06:56:10 buddy vmunix: ip00> Available user memory: 56422 pages, 220.4 Mb
Jun 18 06:56:10 buddy vmunix: ip00> ex[0]: [8:0:14:21:27:39]
```

### 7.1.2.3 Pseudo Terminals

Pseudo terminal master/slave pairs provide connections for remote logins. The script /dev/mknewdevs, by default, makes special files for 32 pseudo terminal pairs at system startup. The master names are /dev/ptyp0 through /dev/ptypf, and /dev/ptyq0 through /dev/ptyqf. The slave names are /dev/ttyp0 through /dev/ttypf, and /dev/ttyq0 through /dev/ttyqf.

The major numbers are 0xc5 (197 – pty) and 0xc4 (196 – tty); the minor numbers, for both pty and tty devices, are 0x0 through 0xf and 0x10 through 0x1f.

Use the /dev/mkptynod command or the /etc/mknod command to add devices at major numbers 0xc5 (197) and 0xc4 (196) if more pseudo terminals are needed. By convention, new master names are /dev/ptyr0 through /dev/ptyrf (minor numbers 0x20 – 0x2f), /dev/ptys0 through /dev/ptysf (minor numbers 0x30 – 0x3f), and so on, and new slave names are /dev/ttyr0 through /dev/ttyrf (minor numbers 0x20 – 0x2f), /dev/ttys0 through /dev/ttysf (minor numbers 0x30 – 0x3f), and so on.

By default, you are restricted to 64 pseudo terminal pairs. Use the /usr/alliant/ptynum command to change the limit. Changes should be in increments of 16, to a maximum of 128.

Each slave pseudo terminal requires an entry in /etc/ttys. The getty process must be specified as none, the terminal type must be specified as network, and the on flag must be omitted or changed to off. The /etc/ttys file as distributed contains entries for ttyp0 through ttypf and ttyq0 through ttyqf.

## 7.1.3 Network Names and Addresses

Names and addresses must be supplied as follows:

- **Hostname** – The hostname for a system is the administrator's invention. The name is a character string; blanks, tabs, newlines, and pound signs (#) are excluded.
- **Network name** – The name of a network is the administrator's invention unless it is a registered Internet name. The name is a character string of up to 12 characters; blanks, tabs, newlines, and pound signs (#) are excluded.
- **Internet address** – The Internet address is the administrator's invention within specifications (see below) unless it is a registered Internet name. The address is a 32-bit entity that is typically specified in a 4-part dot notation.

To obtain information for registering a network on the Internet, send electronic mail to HOSTMASTER@SRI-NIC.ARPA or write to the following address. Registration is required to connect your network to the DARPA Internet or the DDN Internet. Local area networks that you completely control need not be registered.

```
DDN Network Information Center
SRI International
Room EJ217
333 Ravenswood Avenue
Menlo Park, CA 94025
```

An official Internet address identifies a network by class and number, and leaves space for the local administrator to identify the local hosts. Classes are as follows:

- **Class A** – Bit 31 of the address is 0; bits 30:24 identify the network; bits 23:0 are used by the local administrator to identify hosts.
- **Class B** – Bits 31:30 of the address are 10; bits 29:16 identify the network; bits 15:0 are used by the local administrator to identify hosts.
- **Class C** – Bits 31:30 of the address are 11; bits 29:8 identify the network; bits 7:0 are used by the local administrator to identify hosts.

The address can be specified using decimal numbers and dot notation as shown below. Drop the local host portion to identify a network. Include the local host portion to identify a host on a network.

**Network Addressing**

| | | | | |
|---|---|---|---|---|
| class A | 0–127 . 0–255 . 0–255 . 0–255 | | | |
| | network | local host | local host | local host |
| class B | 128–191 . 0–255 . 0–255 . 0–255 | | | |
| | network | network | local host | local host |
| class C | 192–255 . 0–255 . 0–255 . 0–255 | | | |
| | network | network | network | local host |

Registered Internet addresses are class C unless a compelling case is made for a category A or B address. Typically, if you are not using a registered address, you should use an address that follows the specifications for a class C Internet address, that is, the first part of the address is in the range 192–255, the first three parts identify the network, and the last part identifies hosts. This scheme minimizes the update work if, at a later date, you register your network.

For example, if your site has two networks, you might address them as 192.0.1 and 192.0.2. Host addresses for 192.0.1 would be 192.0.1.1, 192.0.1.2, and so forth. Host addresses for 192.0.2 would be 192.0.2.1, 192.0.2.2, and so forth.

Three administration files and one administration command define the names and addresses for a network:

- **/etc/networks** – A file that contains the names and Internet addresses of networks that you might access. The file as distributed contains many entries for a large number of registered Internet networks. You can delete those that you do not want. The file must contain an entry for network number 127, the loopback – or local – network. Add to this file the names and addresses of your local area networks.

- **/etc/hosts** – A file that contains the Internet addresses and names of the local hosts that you might access. The file as distributed contains many entries for a large number of local hosts on registered Internet networks. You can delete those that you do not want. Add to this file the addresses and names of the hosts on your local area networks.

- **/etc/ethers** – A file that contains the Ethernet addresses and host names of systems that you might access. The file is not distributed. Create the file and add to it the Ethernet addresses and names of the systems on your local area networks, one entry per line, separating the address and name by any number of spaces and tabs. A pound sign (#) in column 1 indicates a comment. A name must be the first name specified in an /etc/hosts entry, not an alias.

- **/etc/ifconfig** – A command that configures a network on your system. You must specify the controller (interface) name, the address or name of the network, and various operating parameters. On Alliant systems, **-trailers** must be specified. The start-up script **rc.local** contains a template command to configure your network through Ethernet controller **ex0**. Provide a value in place of **my-netmask** or delete **netmask my-netmask**. Provide a value in place of **my-broadcast** or delete **broadcast my-broadcast**. Change **ex0** if this is not the name of the controller. If you have a second Ethernet controller, add a new **ifconfig** line.

Whenever you change /etc/networks, /etc/hosts, or /etc/ethers on one system, propagate the changes to all other systems in your local networks. If you are running network information service, you should make changes on the master server for each data base and run the **make** command in /etc/yp. If you are not running network information service, edit the files on each system or copy an edited file to each system (for example, with the **rcp** or **ftp** command).

The following example demonstrates the setup for two networks: n1net consisting of the systems n1s1, n1s2, and n1s3; and n2net consisting of the systems n2s1 and n2s2. The files /etc/networks, /etc/hosts, and /etc/ethers must be made identical on all the systems in the networks. The display of the first few lines of /etc/rc.local indicates that the system is n1s1.

```
# cat /etc/networks
# This file has only the essential entries for the sample system.
# The first entry is required for all networks.
loopback-net    127             software-loopback-net
n1net           192.0.1         user-net-1
n2net           192.0.2         user-net-2
# cat /etc/hosts
# Column 1 is the internet address of the host.
# Column 2 is the system name; other columns (optional) are aliases.
# The first line is required for all networks.
127.0.0.1       localhost
192.0.1.1       n1s1    alpha1
192.0.1.2       n1s2    alpha2
192.0.1.3       n1s3    alpha3
192.0.2.1       n2s1    beta1
192.0.2.2       n2s2    beta2
# cat /etc/ethers
# Get the ethernet address from the boot script.
8:0:20:0:A1:61 n1s1
8:0:20:0:A7:35 n1s2
8:0:20:0:C5:63 n1s3
8:0:20:0:C5:6D n2s1
8:0:20:0:C5:CB n2s2
# cat /etc/rc.local
PATH=/etc:/usr/etc:/bin:/usr/ucb:/usr/bin:/usr/alliant

# myname is my symbolic name
# myhostid is my internet address in #.#.#.# format
# my-netmask is specified in /etc/networks
#
hostname n1s1
#domainname mydomain
ifconfig ex0 -trailers broadcast 192.0.0.0 `hostname`
ifconfig lo0 localhost
sleep 3
route add `hostname` localhost 0
hostid `hostname`
...
```

### 7.1.3.1 Setting localhost for Non-networked Machine

If a machine is not included in the network, you must set the *localhost* entry and comment out the *myname* entry in the /etc/hosts file as follows:

```
127.1       localhost       your-host-name
#125.100    myname
```

(The *your-host-name* is excluded when the machine is in a net.)

### 7.1.4 Equivalent (Trusted) Systems And Accounts

The following files permit automatic access (no password is required) to your system from another system:

- **/etc/hosts.equiv** – Lists the names of hosts whose users, except for root, can access your system automatically.
- **/.rhosts** – Lists the names of hosts whose root user can access your system automatically.
- **~/.rhosts** – Lists the names of hosts whose *user* user can access your system automatically, where *user* is the user in whose home directory .rhosts resides. The ~/.rhosts file permits access for selected users if /etc/hosts.equiv is absent or does not contain the user's host.

The user accessing your system must be in /etc/passwd on your system. On equivalent systems, you should copy the password file from one system to another after a change is made. (This procedure is different if you are running Network Information Service.) See the *Concentrix-2800 System Reference Manual: Miscellaneous Routines, Special Files, and File Formats*, hosts.equiv and rhosts for more detailed information on the use of these files.

Specifying systems in /etc/hosts.equiv but not /.rhosts permits user-level access to your system but denies superuser access. Copying /etc/hosts.equiv to /.rhosts permits complete access to your system.

The following example of /etc/hosts.equiv on the system n1s1 permits automatic access to ordinary users from all other systems on the network, and to the superuser (root) from the system n1s2 (anyone logging in as super-user from a system other than n1s2 must supply a password):

```
# cat /etc/hosts.equiv
n1s1
n1s2
n1s3
n2s1
n2s2
# cat /.rhosts
n1s1
n1s2
```

**WARNING:** If the name of an Ethernet terminal server in /etc/hosts is added to /etc/hosts.equiv or /.rhosts, the terminal server may permit unrestricted access to user accounts. If security is an issue, be aware of these problems.

Some terminal servers "fake" the login: prompt and pass the username as part of the rlogin request. If the server is listed in /etc/hosts.equiv, Concentrix-2800 will not require a password to login to the account. Since terminal servers do no authentication, this is a security issue.

Some terminal servers can also pass a null username to rlogind. This will result in a root login (no password required) if a null user/password line exists in /etc/passwd (due to a bug in vipw) and the server's name is in /.rhosts.

## 7.1.5 Examining and Testing the Network

The administration command /etc/ping tests the connection between two systems by sending and receiving packets. After typing the command, you should see a steady display indicating that packets are being sent and received. Type CTRL C to stop the display. The terminating display should indicate 0% packet loss. The following example indicates a good connection between the local host and n1s2:

```
# /etc/ping n1s2
PING n1s2 (125.0.0.3): 56 data bytes
64 bytes from 125.0.0.3: icmp_seq=0. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=1. time=29. ms
64 bytes from 125.0.0.3: icmp_seq=2. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=3. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=4. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=5. time=29. ms
64 bytes from 125.0.0.3: icmp_seq=6. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=7. time=19. ms
64 bytes from 125.0.0.3: icmp_seq=8. time=19. ms
CTRL C
----n1s2 PING Statistics----
9 packets transmitted, 9 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 19/21/29
```

A packet loss greater than 0% indicates a problem. If the other system is down, you do not see a steady display (the command appears to hang). When you type CTRL C, the terminating display indicates 100% packet loss:

```
# /etc/ping n1s2
PING n1s2 (125.0.0.109): 56 data bytes
CTRL C
----n1s2 PING Statistics----
12 packets transmitted, 0 packets received, 100% packet loss
```

The command /usr/ucb/netstat provides status information on your connections to other systems. The command /usr/ucb/netstat -i shows the state of the network interfaces (network controllers and the local software loopback mechanism). The command /usr/ucb/netstat -a shows the state of work in progress over the network.

Use one of the network services to first test the connection to your own system, then to another system. For example, if your system is n1s1, type /usr/ucb/rlogin n1s1 to test the network connection to your own system. Then type /usr/ucb/rlogin n1s2 to test the network connection to another system.

Any of the following error messages indicates a broken connection with the system you are trying to reach:

```
(1)  exn: no carrier
(2)  exn: Ethernet jammed
(3)  Connection timed out
(4)  no server
```

The typical hardware problems are as follows:

- The system you are trying to access is down. Type /usr/ucb/ruptime to get the status of all systems on your local network.
- A cable is loose. Check the machine and transceiver connections.
- A terminator is missing or loose. Make sure all terminators are in place and secure.
- A cable is broken. The cables are somewhat fragile and undue bending or pressure can damage them.

The typical software problems are as follows:

- The system you are trying to access has incorrect entries, or no entries at all, for your system in any or all of the following files: /etc/networks, /etc/hosts, /etc/ethers. Ensure that the entries are present, and that the host name, Ethernet address, and internet address are correct in all instances.
- A network controller is not configured. Ensure that the /etc/ifconfig command (or commands – one per controller) are correct in /etc/rc.local. If you adjust the command(s) in /etc/rc.local, also enter the command(s) manually, or reboot the system.

## 7.1.6 Ethernet Performance

Ethernet performance is impacted by the hardware on the machine, the placement of the softnet portion of the Ethernet driver, and by the amount of space allocated for buffers.

### 7.1.6.1 Placement of Softnet Portion of Ethernet Driver

By default, the softnet portion of the Ethernet driver runs on machine 28 (IP0). You can improve performance by patching the kernel variable soft-net-cpu to place the softnet processing on a low numbered CE that is not in a complex, typically machine 2 or 3.

### 7.1.6.2 Buffers

Running out of buffers for network services seriously degrades network and system performance. However, allocating excessive memory for network buffers can make the system unbootable. By default, Concentrix-2800 allocates as many buffers as is practical based on memory size.

Network services use two types of buffers:

- **mbufs** – Memory buffers. Each mbuf is 256 bytes in size and the allocation is specified in pages (16 per page). You can change the mbuf allocation by patching the kernel variable nmbufs (a longword) and rebooting. The default value is 0 and specifies the system algorithm.

- **hpages** – Half-pages. Each hpage is 2048 bytes in size and the allocation is specified in number of hpages. You can change the hpage allocation by patching the kernel variable ex_n_pages (a word) and rebooting. The default value is 0 and specifies the system algorithm.

The command netstat –m displays the network buffer usage and allocation on your system. The following display shows that the system is using 611 mbufs out of 2048 allocated (note that the display is in number of mbufs, not pages), and 10 half-pages out of 128 allocated.

```
% netstat -m
611/2048 mbufs in use:
        70 mbufs allocated to data
        52 mbufs allocated to packet headers
        183 mbufs allocated to socket structures
        255 mbufs allocated to protocol control blocks
        6 mbufs allocated to routing table entries
        43 mbufs allocated to socket names and addresses
        2 mbufs allocated to interface addresses
10/128 half pages in use (low is 58)
512 Kbytes of mbufs allocated to network (29% in use)
256 Kbytes of half-pages allocated to network (7% in use)
768 Kbytes allocated to network (22% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

If your usage tends to be higher than the default allocation, you might specify the allocation. You should no more than double the allocation of either resource on one try. Be sure to save a copy of /unix before modifying it in case you over-allocate and make the system unbootable. The following example increases the allocation of mbufs to 160 pages (2560 mbufs) and the allocation of half-pages to 160:

```
# adb -w /unix
nmbufs?W a0
_nmbufs:        0x80    =       0xa0
ex_n_pages?w a0
_ex_n_pages:    0x80    =       0xa0
CTRL D
```

**Warning:** Do not apply a patch without first copying the kernel (/unix) and understanding how to boot the system from the saved kernel. (See Appendix B, Patch_kernel and Boot_system, for general instructions on patching the kernel and rebooting.) Apply the patches exactly and pay attention to case; for example, typing W instead of w writes a longword instead of a word. In general, do not copy kernels from one system to another. Be conservative in adjusting values such as mbufs.

## 7.2 UUCP NETWORK

The UUCP network supports file transfers and program execution between systems connected by modems, direct lines (RS-232), and local area networks (Ethernet). Two general commands are supported:

- **uucp** – Unix to unix copy
- **uux** – Unix to unix command execution

The sendmail facility uses the UUCP network for systems not on a local area network.

The following figure illustrates the cable (non-Ethernet) connections over which UUCP works. Direct connections are not typical because (1) local area networks serve the purpose much better and (2) each connection uses up a tty port. Modem connections provide the ability for remote systems to communicate periodically over telephone lines. Not every system needs an outgoing modem; however, a system without an outgoing modem cannot initiate a UUCP session. See Chapter 3 for cabling requirements.

**UUCP Network: via modem (top) and direct (bottom)**

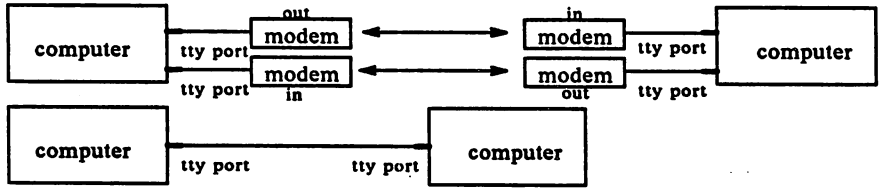


UUCP supports the following user level commands:

| | |
|---|---|
| /bin/rmail | handle remote mail received via UUCP |
| /usr/bin/tip | connect to a remote system |
| /usr/bin/uucp | unix to unix copy |
| /usr/bin/uudecode | decode binary file from mail |
| /usr/bin/uuencode | encode binary file for mail |
| /usr/bin/uulog | display UUCP log files |
| /usr/bin/uuname | list names of UUCP hosts |
| /usr/bin/uuq | examine or manipulate the UUCP queue |
| /usr/bin/uusend | send file to remote host |
| /usr/bin/uux | unix to unix command execution |

UUCP supports the following administration commands:

| | |
|---|---|
| /usr/lib/uucp/uucico | transfer files queued by uucp or uux |
| /usr/lib/uucp/uuclean | UUCP spool directory cleanup |
| /etc/uucpd | UUCP daemon |
| /usr/bin/uupoll | poll a remote UUCP site |
| /usr/bin/uusnap | show snapshot of the UUCP system |
| /usr/bin/uuxqt | UUCP execution file interpreter |

User level commands such as uucp and uux queue jobs to /usr/spool/uucp. The administration command /usr/lib/uucico processes the jobs by making the appropriate remote connections and transferring the jobs to the remote system. The remote system then transfers back any jobs for the local system. A system must have an outgoing modem to initiate job processing. A system with only an incoming modem can only transfer jobs back to another system that calls first.

Proper operation of UUCP requires the existence and correct configuration of the spooling directories /usr/spool/uucp and /usr/spool/uucppublic. You can generate the necessary directories, subdirectories, and files by executing a makefile as follows:

```
# cd /usr/src/usr.bin/uucp
# make mkdirs
```

The following files must be placed in the directory /usr/lib/uucp. Templates and descriptions exist in the directory /usr/src/usr.bin/uucp/UUAIDS; also, see the *Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*. All files must be owned by UUCP.

- **L-devices** – Contains the names of the ports used by UUCP and information on the modems or hardwire connections.
- **L.sys** – Contains calling information on remote systems. The permissions on this file should be set to 400 to prevent users from reading the passwords for the remote systems.
- **USERFILE** – Contains user names and pathname permissions for logging into remote systems. A user name must be present in /etc/passwd on the remote system. A system name must be present in L.sys on the local system.
- **L-dialcodes** – Maps codes in L.sys to full phone numbers.
- **L.aliases** – Provides alternate names for the system names in L.sys.
- **L.cmds** – Lists the commands that can be executed by uux.

The following files should be transferred from /usr/src/usr.bin/uucp/UUAIDS and edited if desired:

- **uu.hourly** – Invokes /usr/lib/uucp/uucico to transfer queued files. Copy this file to /usr/local/lib/uucp.6hours or another file of your choice.
- **uu.daily** – Invokes /usr/lib/uucp/uuclean and performs other housekeeping chores. Copy this file to /usr/local/lib/uucp.daily or another file of your choice.

The two lines in /usr/lib/crontab containing the names of the transferred files should be uncommented. Adjust /usr/lib/crontab if you use different names or change the timing.

```
0 0 * * *       uucp    /usr/local/lib/uucp.daily
7 5,12,18 * * *         uucp    /usr/local/lib/uucp.6hours
```

The directory /usr/src/usr.bin/uucp/UUAIDS contains other files that are unsupported but that might be useful. These files are typically copied to /usr/local or /usr/lib/uucp.

See Add_uucp in Appendix B for detailed procedures.

## 7.3 NETWORK INFORMATION SERVICE

The network information service (NIS) facility permits the distribution of centrally maintained data bases over a local area network. The host systems using one set of data bases are called a domain and are identified by a domain name.

> **Note:** The term "network information service (NIS)" is now used to refer to the service formerly known as Sun Yellow Pages. The functionality remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

The command /bin/domainname prints the name of the domain in which your system participates:

```
# /bin/domainname
d1
```

The NIS data base for a domain consists of one or more maps maintained in the etc/yp/*domainname*. Each map consists of two files with suffixes of .pag and .dir. To list the maps for your domain, use the command /usr/bin/ypwhich -m. The following example shows a display from a domain with the standard maps:

```
# /usr/bin/ypwhich -m
passwd.byname d1
passwd.byuid d1
group.byname d1
group.bygid d1
hosts.byname d1
hosts.byaddr d1
ethers.byname d1
ethers.byaddr d1
networks.byname d1
networks.byaddr d1
services.byname d1
rpc.bynumber d1
protocols.byname d1
protocols.bynumber d1
netgroup d1
netgroup.byuser d1
netgroup.byhost d1
mail.aliases d1
usrclass.byname d1
usrclass.byuid d1
```

To print the contents of a map, apply the ypcat command to the map name. The .byname maps in the standard configuration can be abbreviated to omit the extension. The following example prints the passwd.byname map:

```
# /usr/bin/ypcat passwd | /bin/more
```

The standard maps are derived from the following files: /etc/passwd, /etc/group, /etc/hosts, /etc/ethers, /etc/networks, /etc/services, /etc/rpc, /etc/protocols, /etc/netgroup, /usr/lib/aliases, and /etc/usrclass.

The NIS maps are maintained on one or more hosts designated as NIS servers for the domain. For each map, one of the servers is the master server and the others are slave servers; the simplest arrangement is to use one host as the master server for all maps. Changes to an NIS data base are made to the master server's map and then propagated to any other servers in the domain. A server performs its work through the daemon /etc/ypserv.

A client is a system in the domain that uses data base information. A client performs its work through the daemon /etc/ypbind.

The following example demonstrates a domain with two servers and three clients. One of the servers is the master server for all maps.

**Sample NIS Domain**

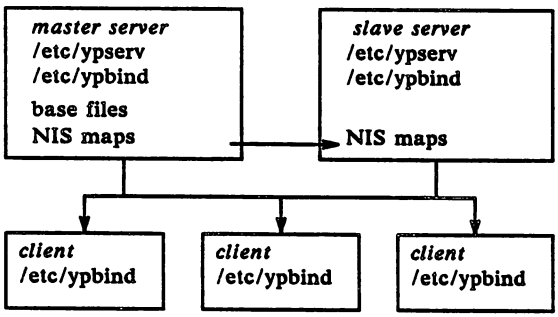


Setting up NIS on a local area network entails the following work:

- Invent a domain name. Decide which system is the master server (you can use different servers as masters for different maps, but one master is simplest), which systems are the slave servers, and which systems are clients only.
- Set the domain name on all participating systems (/bin/domainname).
- Initialize NIS on the master server (/etc/yp/ypinit -m).
- Initialize NIS on the slave servers (/etc/yp/ypinit -s).
- Execute the daemons /etc/ypserv and /etc/ypbind on the servers. Execute the daemon /etc/ypbind on the clients.
- Edit /etc/rc.local on all systems: uncomment the domainname line and specify the correct domain name; uncomment the three ypserv lines (servers only); uncomment the three ypbind lines.

To update the NIS for a domain, update the base files on the master server(s), then run make on the makefile in /etc/yp.

Non-standard maps can be generated by the command /etc/yp/makedbm and propagated to slave servers by the command /etc/yp/yppush.

See Add_yp in Appendix B for detailed procedures.

# APPENDIX A
# ADMINISTRATION COMMANDS

This appendix specifies the administration commands.

# ac

## login accounting

Produces a printout giving connect time for each user who has logged in during the life of the current wtmp file. A total is also produced.

---

**/etc/ac [-w *wtmp*] [-p] [-d] [*person* ...]**

---

### ARGUMENTS

| | |
|---|---|
| -w *wtmp* | specifies an alternate wtmp file (default is /usr/adm/wtmp). |
| -p | prints individual totals; without this option, only totals are printed. |
| -d | produces a printout for each midnight-to-midnight period. |
| *person* | limits the printout to only the specified login names. |

### DESCRIPTION

Init and login maintain the accounting file /usr/adm/wtmp. The file must exist for connect-time accounting to occur. To start accounting, create the file with a length 0. If the file is left undisturbed, it grows without bound, so periodically any information desired should be collected and the file should be truncated.

### FILES

| | |
|---|---|
| /usr/adm/wtmp | connect-time accounting file |

### SEE ALSO

accton, init, sa
*Concentrix-2800 Commands and Applications Manual*: login
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: utmp

# accton

## enable/disable system accounting

Enables or disables system accounting depending of the presence or absence or an argument.

The accton command is typically invoked by rc.

---

**/etc/accton** [*file*]

---

### ARGUMENTS

*file* places system accounting information for every process executed at the end of the specified file. If omitted, accounting is turned off.

### FILES

| | |
|---|---|
| /usr/adm/acct | execution accounting file |
| /usr/adm/savacct | summary of execution accounting file |
| /usr/adm/usracct | per-user summary of execution accounting file |

### SEE ALSO

ac, sa
*Concentrix-2800 System Reference: System Calls and C Library Routines*: acct

# arcv

## convert archives to new format (old)

(Available but not supported.) Converts archive files from 32v and Third Berkeley editions to a new portable format. The conversion is done in place, and the command refuses to alter a file not in old archive format.

---

**/usr/ucb/arcv** *file ...*

---

### DESCRIPTION

Old archives are marked with the magic number `0177545` at the start; new archives have a first line `!<arch>`.

### FILES

| | |
|---|---|
| /tmp/v* | temporary copy |

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: ar
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ar

# arp

## address resolution display and control

Displays and modifies the Internet-to-Ethernet address translation tables used by the address resolution protocol.

---

**/etc/arp** *hostname*

**/etc/arp -a** [*vmunix* [*kmem*]]

**/etc/arp -d** *hostname*

**/etc/arp -s** *hostname ether-addr* [temp] [pub] [trail]

**/etc/arp -e** *filename*

**/etc/arp -f** *filename*

---

### ARGUMENTS

*hostname* by itself, displays the current ARP entry for *hostname*. The host can be specified by name or by number using Internet dot notation.

**-a** displays all of the current ARP entries by reading the table from the file *kmem* (default /dev/kmem) based on the kernel file *vmunix* (default /vmunix).

**-d** deletes an entry for *hostname*; requires superuser privilege.

**-e** reads *filename* and sets multiple entries in the ARP tables. Entries in the file have the following form with argument meanings as specified above (same as -f with first two arguments reversed; suitable for network information service /etc/ethers):

*ether-addr hostname* [temp] [pub]

**-f** reads *filename* and sets multiple entries in the ARP tables. Entries in the file have the following form with argument meanings as specified above:

*hostname ether-addr* [temp] [pub]

**-s** creates an ARP entry for *hostname* with the Ethernet address *ether-addr*. The Ethernet address is specified as six hexadecimal bytes separated by colons. The entry is permanent unless the keyword temp is specified in the command. If the keyword pub is specified, the entry is published, for example, the system will respond to ARP requests for *hostname* even though the hostname is not its own. The keyword trail indicates that trailer encapsulations may be sent to this host.

## SEE ALSO

ifconfig
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: inet
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: arp

# automount

## automatically mount NFS file systems

Automount is a daemon that automatically and transparently mounts an NFS file system as needed.

---

**automount** [ **-mnTv** ] [ **-D** *name= value* ] [ **-f** *master–file* ]
[**-M** *mount–directory* ] [ **-tl** *duration* ] [ **-tm** *interval* ]
[ **-tw** *interval* ] [ *directory map* [ *-mount–options* ] ] ...

---

### ARGUMENTS

**-m** Suppress initialization of directory-map pairs listed in the auto.master NIS database.

**-n** Disable dynamic mounts. With this option, references through the automount daemon only succeed when the target filesystem has been previously mounted. This can be used to prevent NFS servers from cross-mounting each other.

**-T** Trace. Expand each NFS call and display it on the standard output.

**-v** Verbose. Log status messages to the console.

**-D** *envar=value*
Assign value to the indicated automount (environment) variable.

**-f** *master–file*
Read a local file for initialization, ahead of the auto.master NIS map.

**-M** *mount-directory*
Mount temporary file systems in the named directory, instead of /tmp_mnt.

**-tl** *duration*
Specify a duration, in seconds, that a file system is to remain mounted when not in use. The default is 5 minutes.

**-tm** *interval*
Specify an interval, in seconds, between attempts to mount a filesystem. The default is 30 seconds.

**-tw** *interval*
Specify an interval, in seconds, between attempts to unmount filesystems that have exceeded their cached times. The default is 1 minute.

### DESCRIPTION

Automount is a daemon that automatically and transparently mounts an NFS file system as needed. It monitors attempts to access directories associated with an automount map, along with any directories or files that reside under them. When a file is to be accessed, the daemon mounts the appropriate NFS file system. You can assign a map to a directory using an entry in a direct automount map, or by specifying an indirect map on the command line.

The automount daemon appears to be an NFS server to the kernel. automount uses the map to locate an appropriate NFS file server, exported file system, and mount options. It then mounts the file system in a temporary location, and replaces the file system entry for the directory or subdirectory with a symbolic link to the temporary location. If the file system is not accessed within an appropriate interval (five minutes by default), the daemon unmounts the file system and removes the symbolic link. If the indicated directory has not already been created, the daemon creates it, and then removes it upon exiting.

Since the name-to-location binding is dynamic, updates to an automount map are transparent to the user. This obviates the need to "pre-mount" shared file systems for applications that have "hard coded" references to files.

If you specify the dummy directory /–, automount treats the *map* argument that follows as the name of a direct map. In a direct map, each entry associates the full pathname of a mount point with a remote file system to mount.

If the directory argument is a pathname, the *map* argument points to an indirect map. An indirect map, contains a list of the subdirectories contained within the indicated directory. With an indirect map, it is these subdirectories that are mounted automatically.

The *–mount–options* argument, when supplied, is a comma-separated list of mount options, preceded by a –. If mount options are specified in the indicated map, however, those in the map take precedence.

### Environment

Environment variables can be used within an automount map. For instance, if $HOME appears within a map, automount expands it to its current value for the HOME variable.

If a reference needs to be protected from affixed characters, you can surround the variable name with curly braces.

## USAGE

### Map Entry Format

A simple map entry (mapping) takes the form:

```
directory [ –mount–options ] location ...
```

where *directory* is the full pathname of the directory to mount when used in a direct map, or the basename of a subdirectory in an indirect map. *mount–options* is a comma-separated list of mount options, and *location* specifies a remote filesystem from which the directory may be mounted. In the simple case, *location* takes the form:

```
host:pathname
```

Multiple *location* fields can be specified, in which case automount sends multiple mount requests; automount mounts the file system from the first host that replies to the mount request. This request is first made to the local net or subnet. If there is no response, any connected server may respond.

If *location* is specified in the form:

```
host:path:subdir
```

*host* is the name of the host from which to mount the file system, *path* is the pathname of the directory to mount, and *subdir*, when supplied, is the name of a subdirectory to which the symbolic link is made. This can be used to prevent duplicate mounts when multiple directories in the same remote file system may be accessed. With a map for /home such as:

```
able homeboy:/home/homeboy:able
baker       homeboy:/home/homeboy:baker
```

and a user attempting to access a file in /home/able, automount mounts homeboy:/home/homeboy, but creates a symbolic link called /home/able to the able subdirectory in the temporarily-mounted filesystem. If a user immediately tries to access a file in /home/baker, automount needs only to create a symbolic link that points to the baker subdirectory; /home/homeboy is already mounted. With the following map:

```
able homeboy:/home/homeboy/able
baker       homeboy:/home/homeboy/baker
```

automount would have to mount the filesystem twice.

A mapping can be continued across input lines by escaping the NEWLINE with a backslash. Comments begin with a # and end at the subsequent NEWLINE.

### Directory Pattern Matching

The '&' character is expanded to the value of the *directory* field for the entry in which it occurs. In this case:

```
able homeboy:/home/homeboy:&
```

the & expands to able.

The '*' character, when supplied as the *directory* field, is recognized as the catch-all entry. Such an entry resolves to any entry not previously matched. For instance, if the following entry appeared in the indirect map for /home:

```
*       &:/home/&
```

this allows automatic mounts in /home of any remote file system whose location can be specified as:

```
hostname:/home/hostname
```

### Hierarchical Mappings

A hierarchical mapping takes the form:

```
directory [ /[subdirectory [ -mount-options ] location ... ] ...
```

The initial / within the '/[subdirectory]' is required; the optional subdirectory is taken as a filename relative to the directory. If *subdirectory* is omitted in the first occurrence, the / refers to the directory itself.

Given the direct map entry:

```
/usr/local   \
             /      -ro,intr      loco:/usr/local        alt:/usr/local    \
             /bin   -ro,intr      alt:/usr/local/bin     loco:/usr/local/bin \
             /man   -ro,intr      loco:/usr/local/man    alt:/usr/local/man
```

automount automatically mounts /usr/local, /usr/local/bin, and /usr/local/man, as needed, from either loco or alt, whichever host responded first.

### Direct Maps

A direct map contains mappings for any number of directories. Each directory listed in the map is automatically mounted as needed. The direct map as a whole is not assocated with any single directory.

### Indirect Maps

An indirect map allows you to specify mappings for the subdirectories you wish to mount under the directory indicated on the command line. It also obscures local subdirectories for which no mapping is specified. In an indirect map, each directory field consists of the basename of a subdirectory to be mounted as needed.

### Included Maps

The contents of another map can be included within a map with an entry of the form:

```
+mapname
```

mapname can either be a filename, or the name of an NIS map, or one of the special maps described below.

### Special Maps

There are three special maps currently available: -hosts, -passwd and -null. The -hosts map uses the NIS hosts.byname map to locate a remote host when the hostname is specified. This map specifies mounts of all exported file systems from any host. For instance, if the following automount command is already in effect:

```
automount /net -hosts
```

then a reference to /net/hermes/usr would initiate an automatic mount of all file systems from hermes that automount can mount; references to a directory under /net/hermes will refer to the corresponding directory on hermes. The -passwd map uses the passwd database to attempt to locate the home directory of a user. For instance, if the following automount command is already in effect:

```
automount /homes -passwd
```

then if the home directory for a user has the form /dir/server/username, and server matches the host system on which that directory resides, automount mounts the user's home directory as: /homes/username.

For this map, the tilde character '~' is recognized as a synonym for the username.

The -null map, when indicated on the command line, cancels a previous map for the directory indicated. It can be used to cancel a map given in auto.master.

### Configuration and the auto.master map

automount normally consults the auto.master NIS configuration map for a list of initial automount maps, and sets up automatic mounts for them in addition to those given on the command line. If there are duplications, the command line arguments take precedence. This configuration database contains arguments to the automount command, rather than mappings; unless -f is in effect, automount does not look for an auto.master file on the local host.

Maps given on the command line, or those given in a local auto.master file specified with -f override those in the NIS auto.master map. For instance, given the command:

```
automount -f /etc/auto.master /home -null /- /etc/auto.direct
```

and a file named /etc/auto.master that contains:

```
/homes -passwd
```

automount mounts home directories under /homes instead of /home, in addition to the various directories specified in the /etc/auto.direct file.

## NOTES

- When it receives signal number 1, SIGHUP, automount rereads the /etc/mtab file to update its internal record of currently mounted file systems. If a file system mounted with automount is unmounted by a umount command, automount should be forced to reread the file.
- Shell filename expansion does not apply to objects not currently mounted.
- Since automount is single-threaded, any request delayed by a slow or non-responding NFS server delays all subsequent automatic mount requests until it completes.
- Programs that read /etc/mtab and then touch files that reside under automatic mount points introduce further entries to the file.
- Automatically mounted file systems are mounted with type *ignore*; they do not appear in the output of either mount or df.

## FILES

| | |
|---|---|
| /tmp_mnt | directory under which filesystems are dynamically mounted |

## SEE ALSO

mount
*Concentrix-2800 Commands and Applications*: df
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: passwd

# batchd

## batch daemon

Reads the file /etc/qconf to determine the current configuration of the batch queueing system and maintains status files that are used by the program /usr/lib/batch/qstat. If the daemon receives a signal 15 (SIGTERM), it attempts an orderly shutdown of all the queueing system processes by sending a SIGTERM to all of its own children.

---

## /usr/lib/batch/batchd

---

### DESCRIPTION

The batch system is controlled by the batch daemon, which is started by /etc/init through /etc/rc and wakes up every few seconds to perform batch system services. In particular, batchd is respsonsible for assigning requests to devices and keeping the batch job stream moving. Batchd first checks for new submit and delete requests, performing appropriate device flushes and queue-to-device mappings as required. Then, batchd looks for delayed requests that have reached their activation times and places them in the appropriate queues. Finally, batchd assigns requests to idle devices as determined by queue to device mappings. Batchd also maintains several informational files that are used by the status program qstat. Since only batchd knows the actual ordering of requests in the queues, it writes out a compacted list of the queue requests in the order that batchd has them in core. This information is used by qstat to give an accurate indication of the queueing order.

Three programs provide batch network services. Netsend sends batch requests to remote machines via error-free network connections. The requests are received on the remote system by netdae, a daemon that waits for incoming batch TCP connections on a specified socket and submits the incoming requests to netrecv, which enters the requests into the local queueing system. Netrecv expects a steady input stream of data after which it prints a message indicating the success or failure of the transfer. A message indicating success serves as a transfer of responsibility from the originating site to the receiving site. Netrecv does little data verification; it relies on having a error-free network connection such as a TCP/IP virtual circuit.

The batch system spooling hierarchy allows only batch system programs to enter requests into the queueing directories. The spooling directory hierarchy is as follows:

```
/usr/spool/batch/                 mode 755
/usr/spool/batch/qtmp/            mode 777
/usr/spool/batch/lock/            mode 700
/usr/spool/batch/lock/home/       mode 755
/usr/spool/batch/lock/cntrl/      mode 777
/usr/spool/batch/lock/data/       mode 777
/usr/spool/batch/lock/new/        mode 777
/usr/spool/batch/lock/mod/        mode 777
/usr/spool/batch/lock/adm/        mode 777
/usr/spool/batch/lock/hold/       mode 777
```

The working directory of batchd and other batch processes while they are running is usually the cntrl directory, where all the request files are kept. The directory data contains the shell command files corresponding to the batch requests. The new directory is normally empty and is used to signal batchd that a new request is to be processed. The adm directory contains batchd's queue status files, lock files, and device status files. The hold directory is used by batchd to save copies of the request files for those requests that cause uncorrectable or severe errors to batchd or other batch system processes.

If any jobs are no longer desired when the batch system restarts after a crash, use the qdel -k command to kill them.

## FILES

| | |
|---|---|
| /usr/spool/batch | batch spooling directory |
| /etc/qconf | batch system configuration file |

## SEE ALSO

netdae, netrecv, netsend
*Concentrix-2800 Commands and Applications Manual*: qsub, qdel, qstat
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: qconf

# biod

## NFS daemon

The biod command starts asynchronous block I/O daemons that are used on an NFS client to handle read-ahead and write-behind of buffer cached blocks.

---

**/etc/biod** [*nservers*]

---

### ARGUMENTS

*nservers* is the number of asynchronous block I/O daemons to start. A good value for *nservers* is four; the default is one.

### SEE ALSO

mountd, nfsd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: exports

# bootparamd

## boot parameter server

Provides information to diskless clients necessary for booting. Consults either (1) the NIS boot parameters data base if the NIS service is running, or (2) the file /etc/bootparams if the NIS service is not running.

This command can be invoked through inetd or by the user.

---

**/usr/etc/rpc.bootparamd [-d]**

---

### OPTIONS

**-d** displays debugging information.

### FILES

/etc/bootparams local boot parameters data base

### SEE ALSO

inetd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: bootparams

# catman

## create on-line documentation output directory

Executes nroff -h -mung on files in man directories and stores the output in files with the same names in corresponding cat directories. If necessary, catman creates the cat directories.

---

**/etc/catman** [*option* ...] [*list*]

---

### ARGUMENTS

**-M** *path*[:*path*]...
: specifies the top of one or more man directory hierarchies, for example, /usr/man or /usr/local/man:/usr/man. The default is -M man.

  You can also specify the man directories prior to executing catman by setting the environment variable MANPATH to *path*[:*path*]... as in -M. The -M option overrides any MANPATH specification.

**-n**
: prevents creation of the whatis database.

**-p**
: prints out what would be done instead of doing it.

**-v**
: (verbose) prints the name of each file created.

**-w**
: only creates the whatis database.

*list*
: specifies a list of man sections. For example, 12f refers to man1, man2, and manf. The default is all sections.

### DESCRIPTION

The man directory hierarchy consists of subdirectories named man*n* and cat*n*, where *n* is a digit or a letter. Each file in each directory specified by *list* is examined and those files whose formatted (cat) versions are missing or out of date are processed.

### WHATIS DATABASE

Unless -n is specified, catman invokes /usr/lib/makewhatis to create a new whatis data base in a file named whatis in the current working directory. Note that if you specify more than one man data base in -M or PATHNAME, the earlier whatis data bases are overwritten. The whatis data base should be moved to the top directory of the man directory hierarchy, for example, /usr/man/whatis.

You can explicitly invoke the script /usr/lib/makewhatis to reconstruct the whatis data base. The script takes a single optional argument, the pathname of the man data base, which defaults to /usr/man. The man subdirectories that are processed are fixed in the script at 0 through 8, a, c, f, r, and s. The script employs the services of the program /usr/lib/getNAME to extract the whatis information from the man files.

## EXAMPLES

The following example builds cat files, as necessary, for all sections of the man data base and rebuilds /usr/man/whatis:

```
# cd /usr/man
# /etc/catman
```

The following example builds cat files for section 1 of the man data base that starts at /usr/proto/man. The whatis data base is not built.

```
# /etc/catman -n -M /usr/proto/man 1
```

## FILES

| | |
|---|---|
| /usr/lib/getNAME | program that extracts whatis line from man files |
| /usr/man/whatis | whatis data base |
| /usr/man/man?/* | on-line documentation in nroff source format |
| /usr/man/cat?/* | on-line documentation in nroff output format |

## SEE ALSO

makewhatis
*Concentrix-2800 Commands and Applications Manual*: apropos, man, nroff, whatis, whatisin

# chown

## change owner

Changes the owner of specified files. Only the super-user can change owner, in order to simplify accounting procedures.

---

**/etc/chown [-f] [-R]** *owner*[*.group*] *file ...*

---

### ARGUMENTS

*owner* is either a numeric user identification (uid) or a login name, as specified in the password file (/etc/passwd).

*group* is either a numeric group identification (gid) or a group name, as specified in the group file (/etc/group).

**-f** disables error reporting.

**-R** recursively changes the ownership of files under any directories encountered. The ownership of symbolic links are changed but the links are not traversed.

### EXAMPLES

The following example changes the owner of all the files in the current directory to hacker:

```
# /etc/chown hacker *
```

The following example changes the owner of all the files in the current directory and, recursively, all subdirectories to hacker and the group to algorithms:

```
# /etc/chown -R hacker.algorithms *
```

### FILES

/etc/passwd

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: chgrp
*Concentrix-2800 System Reference: System Calls and C Library Routines*: chown
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: passwd, group

# clinfo

## print cluster information

Prints a message listing the number of IPs and CEs, the number of processors detached and attached, and the composition of each cluster.

---

## /usr/alliant/clinfo

---

### EXAMPLE

```
# clinfo
    Number of IPs =  1
    Number of CEs = 12
        Detach/Attach = 5/8
            Cluster 0: CE00 CE01 CE04 CE05
            Cluster 1: CE08 CE09
            Cluster 2: CE10 CE11
```

### SEE ALSO

clcr, clrm

*Concentrix-2800 System Reference: System Calls and C Library Routines*: cluster_create, cluster_info, cluster_remove

# clri

## clear i-node

Writes zeros to the specified i-nodes on the specified filesystem. Any blocks in the affected file are listed as missing in a subsequent file system check.

NOTE: Use fsck for normal file system repairs, not clri.

---

## /etc/clri *filesystem i-node*

---

### ARGUMENTS

*filesystem* name of the file system. The name can be the mount point, the raw device name, or the block device name. Defaults to all the file systems in /etc/fstab.

*i-node* number of an i-node.

### DESCRIPTION

The primary purpose of this routine is to remove a file that appears in no directory. If you clear an i-node that does appear in a directory, also remove the directory entry. Otherwise, when the i-node is reallocated to a new file, the old directory entry points to the new file. At that point, removing the old entry destroys the new file and the new entry again points to an unallocated i-node.

### NOTES

- The file should be closed.

### SEE ALSO

fsck

# clrm

## remove cluster

Deletes a specified computational cluster; all processors in that cluster become detached. You must have superuser privilege.

---

## /usr/alliant/clrm *clusternumber*

---

### ARGUMENTS

*clusternumber* specifies the cluster to delete.

### DIAGNOSTICS

- `Create failed`: Not owner – Either the feature is disabled or you are not superuser.
- `Illegal complex number` *foo* – The argument is not a number.

### EXAMPLE

```
# clrm 1
```

### SEE ALSO

clinfo, clrm

*Concentrix-2800 System Reference: System Calls and C Library Routines*: cluster_create, cluster_info, cluster_remove

# compile

## compile the terminfo database

Compile translates the source files in the terminfo terminal capability database into their object format.

---

**/usr/lib/terminfo/compile [-v [*n*]] *file***

---

The actual location of the compile command is determined by the system administrator.

### ARGUMENTS

**-v** obtains debugging and tracing information.

*n* controls the amount of debugging information given, according to the following table:

| | |
|---|---|
| 1 | Names of files created and linked (the default) |
| 2 | Information related to the "use" facility |
| 3 | Statistics from the hashing algorithm |
| 5 | String-table memory allocations |
| 7 | Entries into the string-table |
| 8 | List of tokens encountered by scanner |
| 9 | All values computed in construction of the hash table |

*file* is a file that contains one or more terminfo entries, as described in terminfo (*File Formats*). This file need not be self-contained; it may contain "use" entries which refer to terminals not described fully in the same file as long as those entries have been compiled at some earlier time.

### DESCRIPTION

The object files are normally placed in subdirectories of the directory /usr/lib/term (see the file format term), but if the environment variable TERMINFO is defined, it is taken to be the name of an alternate directory to use.

### FILES

/usr/lib/term/* Default location of object files

### SEE ALSO

dumpterm
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: terminfo, term

# comsat

## biff server

Listens for reports of incoming mail and notifies users if they have requested this service.

The comsat daemon is started as needed by /etc/inetd.

---

## /etc/comsat

---

### DESCRIPTION

Comsat listens on a datagram port associated with the biff service specification for one-line messages of the form:

*user@mailbox-offset*

If the specified user is logged into the system and the user's terminal has the owner execute bit turned on (set by the command biff y), *offset* is used as a seek offset into the appropriate mailbox file and the first 7 lines or 560 characters of the message are printed on the user's terminal. Lines that appear to be part of the message header other than the "From," "To," "Date," and "Subject" lines are not included in the displayed message.

### EXAMPLES

Not intended for user invocation.

### FILES

| | |
|---|---|
| /etc/utmp | login record of current users |

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: biff
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services

# create_root

## generate root partition for diskless client

Customizes the prototype diskless client root tree for a particular client. The customization includes copying the prototype tree into the client's root path, making the appropriate devices in the client's /dev directory, and customizing the administration files in the client's /etc directory.

---

**/usr/etc/install/create_root** *host yptype root home exec share server domain*

---

### ARGUMENTS

Running the command with no arguments lists the available options.

*host* is the name of the client whose root you are creating.

*yptype* is the NIS mode of operation for the client; one of the following: master, slave, client, or none. Typical values are client or none.

*root* is the full pathname on the server where the client's root tree is being installed, for example, /export/root/client1.

*home* is the path being used for user home directories and is typically /home. If another server is providing this service, the path can have the server name and colon prepended, as in servername:/home. Can also be specified as none if this featue is not used.

*exec* is the full pathname on the server where the executable files for the client's architecture are located, for example, /export/exec/sun3.

*share* is the full pathname on the server where the sharable files are located, for example, /export/exec/share. The client is set up to mount this path on its local /usr/share directory. If this mount is not desired, specify none instead of a pathname.

*server* is the host name of the server machine.

*domain* is the NIS domain name for the client.

### SEE ALSO

INSTALL, setup_client

# cron

## clock daemon

Executes commands at specified dates and times according to the instructions in the files /usr/lib/crontab and /usr/lib/crontab.local; either, neither, or both files may exist.

Cron never exits and should be invoked only once. Typically, cron is invoked by rc.

---

## /etc/cron [-d *logfile*]

---

### ARGUMENTS

-d *logfile* enters debugging information into *logfile*.

### DESCRIPTION

A crontab file consists of lines of seven fields each. The fields are separated by spaces or tabs, except that the last field consists of all text following the sixth field. The first five fields are integer patterns that specify:

- minute (0–59)
- hour (0–23)
- day of the month (1–31)
- month of the year (1–12)
- day of the week (1–7 where 1 = Monday)

Each of these patterns can contain: a number in the range above; two numbers separated by a dash to mean an inclusive range; a list of numbers separated by commas to mean any number in the list; or an asterisk (*) to mean all legal values.

The sixth field is a user name. The command that follows is executed with the specified user's user identification and permissions.

The seventh field is a command that is executed by the shell at the specified times. A per cent character in this field is translated to a new-line character. Only the first line (up to a percent sign or end of line) of the command field is executed by the shell. The other lines are made available to the command as standard input.

Cron checks by crontab's modification date every minute, and includes the changes if the modification date changed.

## NOTES

Users of cron should be aware of the following information:

- cron forks and then execs /bin/sh to execute its commands, effectively performing

  ```
  /bin/sh -c cmd
  ```

  *cmd* will be performed in the context of the new shell.

- commands executed by cron inherit cron's environment.

  cron executes commands by forking and then exec'ing a shell to execute the command. The exec'd shell inherits the environment of cron. This may affect the operation and results of commands submitted to cron which depend on environment variables.

- If *cmd* is a script of shell commands, it is only considered executable by the kernel if the first line is of the form

  ```
  #! [ <ws> ] shell_interpreter [ <ws> shell_arguments ] [ <ws> ] <newline>
  ```

  where <ws> means whitespace. The kernel runs the shell_interpreter using the shell_arguments on the rest of the file.

  If not of this form, then the script file is further examined by the shell (sh or csh) to determine how to handle it.

- the csh may terminate execution of the script prematurely if it encounters a reference to an undefined shell variable.

## FILES

/usr/lib/crontab

## SEE ALSO

exec, csh, sh

# dfmon

## dynamic disk summary and usage

The dfmon command gives a continuous display of non-NFS mounted disks, showing the number of bytes allocated to each partition, the number used, the number free, and the percentage used; inodes are treated similarly.

---

**/usr/alliant/dfmon [-i#] [-d#] ] [-n#] [-s{dfmp}]**

---

### ARGUMENTS

**-i#** pause for interval of # seconds between displays of the current disk statistics. Default is 10 seconds.

**-d#** highlight any disk whose filesystem usage is above #. Default is 90%.

**-n#** highlight any disk whose node usage is above #. Default is 75%.

**-s** sort the display by:

| | |
|---|---|
| **d** | disk name |
| **f** | free space |
| **m** | mount point |
| **p** | percent used |

Default is **-sm**.

### DESCRIPTION

You can change the arguments of dfmon while it is running.

The "diff" column gives the difference in available disk space between the current and last sample. A negative number indicates disk space consumed, a positive number indicates space released.

Type CTRL C or **q** at any point to quit dfmon.

### EXAMPLE

The following example shows a typical dfmon display.

```
columbo% dfmon -i1
Wait is 1 seconds, usage_limit 90%, inode_limit 75%,  Wed Jun 19 16:58:42 1991
 Disk      Mounted     kbytes    used    free   avail    diff used  iused  ifree used
/xdv00r   /           37.2M   27.1M   10.0M    6.3M   0.0K   81%   1.1K   5.0K  17%
 /stripe1 /dumps     698.3M  502.8M  195.5M  125.6M   0.0K   80%   0.0K 143.3K   0%
 /xdv31v  /graf      121.7M  100.1M   21.6M   18.0M   0.0K   84%   0.4K  52.9K   0%
 /stripe0 /t0        493.4M  427.2M   66.2M   46.5M   0.0K   90%   4.3K 121.2K   3%
 /xdv01h  /u0        770.4M  603.5M  167.0M   90.0M   0.0K   87%  17.7K  78.6K  18%
 /xdv21h  /u1        243.1M  169.8M   73.3M   49.0M   0.0K   77%   2.7K 103.8K   2%
 /xdv31u  /u2        121.4M   68.5M   52.9M   49.2M   0.0K   58%   1.7K  51.6K   3%
 /xdv00f  /usr       354.0M  292.2M   61.8M   26.4M   0.0K   91%  16.6K  28.5K  36%
```

# diskf

## disk utility program

Performs various tasks on disks such as formatting, verifying ecc (error correcting code), exercising, bad block handling, and remapping bad tracks.

---

**/etc/diskf** [/dev/*raw_device*] *command* [*from* [*to*]] [*flawfile*]

---

### ARGUMENTS

*raw-device* specifies the raw device name of a disk partition, for example, rxdv01z.

You can set the environment variable DISKF_SCRATCH_DISK to the name of the special file you wish to work on. If a special file name is specified as argument one of diskf, it overrides the environment variable. Otherwise, diskf uses the value of the environment variable.

*command* specifies a diskf command, optionally followed by one or more modifiers, as described below; each command and modifier is a single character; no spaces must separate a command and its modifiers.

*from* [*to*] if *from* and *to* are omitted, the command applies to the entire partition specifed by the raw device name (argument 1). If *from* is specified alone, the command applies to a unit within the partition specifed by the raw device name. If *from* and *to* are both specified, the command applies to a range of units within the partition specifed by the raw device name. Specify *from* and *to* as follows:

*block*
*cylinder.*
*cylinder.head*
*cylinder.head.sector*
end (applies only to *to*)

*block* a decimal number that specifies a disk block within the partition.

*cylinder* a decimal number that specifies a cylinder within the partition.

*head* a decimal number that specifies a head. Defaults to the first head (head 0) for *from* and the last head (head *n*–1, where *n* is the number of heads) for *to*.

*sector* a decimal number that specifies a sector. Defaults to the first sector (sector 0) for *from* and the last sector (sector *n*–1, where *n* is the number of sectors) for *to*.

end (applies only to *to*) the last cylinder head and sector for the partition.

NOTE: On SCSI disks, block is the only unit guaranteed to be accurate.

Some valid specifications:

| | |
|---|---|
| 100 | block number 100 |
| 10. | cylinder 10 |
| 10.5 | cylinder 10, head 5 |
| 10.5.1 | cylinder 10, head 5, sector 1 |

*flaw-file* can be specified with the aF, aZ, B, f, F, D and C commands. Typically, the flaw file is named *device-type_serial-number*. The flaw file contains a list of media defects, one per line, in the form:

*cylinder head bytes-from-index bit-length*

An internal flaw map, containing the same information as the corresponding flaw file, can be maintained on the last track of the first cylinder of each disk and is specified by the keyword internal.

The internal flaw map is not maintained on SCSI drives. diskf f is the only command that uses the *flaw-file* on SCSI disks; it operates in a slightly different manner. See the discussion of SCSI disks for the p, pp, s, and f commands.

## COMMANDS

**NOTE:** Commands a*, b, B, C, D, F, and P do not apply to SCSI disks; f and s work in a different manner on these disks.

**aF** (alternate format) (SMD and IPI-2 disks) all tracks within the range having too many flaws per track are mapped out to the t partition. Do this after formatting the disk. Option aFp attempts to preserve data.

**am** (alternate map) (SMD and IPI-2 disks) maps a specific track out to the t partition. Option amp attempts to preserve data.

**ap** (alternate print) (SMD and IPI-2 disks) prints currently mapped out tracks.

**au** (alternate unmap) (SMD and IPI-2 disks) undoes mapping of a selected track. If the specified track is not mapped, an error message is received.

**aZ** (alternate zero) (SMD and IPI-2 disks) initializes the alternate track area. Do this after formatting the drive. Uses either the internal flaw map, or one specified by the user. The flaw map is needed because of the potential for bad tracks in the alternate track area. Tracks with flaws in the t area are currently marked bad and not used.

**b** (boo-boos) (SMD and IPI-2 disks) scans the disk and prints the sectors that were slipped to standard output in flaw file format; useful if the flaw file is lost.

**B** (bad spots) (SMD and IPI-2 disks) prints the internal flaw map to standard output in flaw file format.

**C** (compare) (SMD and IPI-2 disks) scans the disk and compares the sectors that were slipped to the specified flaw file or (if no flaw file is specified) to the internal flaw map. Extra slips in either the flaw file (or internal flaw map) or on the disk are printed to standard output in flaw file format. This command can be used to ensure that the f command ran properly.

**D** (defects to disk) (SMD and IPI-2 disks) writes the specified flaw file to the internal flaw map on the disk. The old internal flaw map is written out to a file named *special-file* old_flaws where *special-file* is the last segment of the pathname used as the special file.

**f** (format) formats the disk. If a flaw file is specified, the bad sectors specified in the flaw file are slipped during the formatting process.

On a SCSI disk, the entire SCSI drive is formatted. For this reason, use the f command with either no arguments or the arguments 0 end. Use of any other argument yields an error. When you format the drive with the f command, the drive reads its own PLIST (the vendor's flaws), marks bad spots, and formats around them.

It is possible to provide additional flaws which you wish to have mapped out. It is usually desirable to provide the "grown" defect list, or GLIST. To do this, specify a *flaw-file* in its usual place. See p, pp, and s below.

**F** (format from file) (SMD and IPI-2 disks) formats only the tracks specified in the flaw file, slipping the bad sectors. To completely format disks, the f command with flaw file is recommended.

**h H** (help) prints on-line help.

**i** writes /etc/booter to its proper place on the disk, taking into account block size.

**L** (label) reads the label from the partition and prints status information as in the S command. The L command does not convert automatically to the z partition as S does.

**M** obsolete - use the F command.

**O** (on-line) brings a disk that was not ready during system installation on-line.

**p** (print flaws) reads the hardware flaw map and prints the flaws for the *entire drive* to standard output, no matter what partition you specify.

**Note:** The t partition on fj1000s disks contains the alternate track area. You should not access the t partition, except through the alternate track commands of diskf, and you should not create a special file (/dev_entry) for a t partition. The t partition always gives errors if read; its headers look wrong to a controller trying to read them directly.

On an SMD disk, typically, you redirect this output to a file named *device-type_serial-number*, for example, fj690s_87053.

**Warning:** You must read the hardware flaw map before the disk is formatted. Once a disk is formatted, this hardware flaw information is lost. After formatting, the flaw file should be written to the internal flaw map using the D command. Store a copy of the flaw file in a safe place just in case the internal flaw map is lost.

If the utility is reading a SCSI disk, the command prints out the list of "grown flaws" known as the GLIST. You can redirect the output of the p command to a file, then use that file as input for the format command.

On an IPI-2 disk, the command prints out the vendor's flaws *with* parallel head geometry correction.

**pp** (IPI-2 and SCSI disks) reads the vendor's list of flaws for the entire drive and prints them to standard output. This is the raw vendor flaw map.

On an IPI-2 disk, the command prints out the vendor's flaws *without* parallel head geometry correction.

**P** (print flaws) (SMD drives) reads the hardware flaw map over a range and prints it to standard output in flaw file format. In general, use **p**.

**r** (read) reads the data from the disk and prints it on standard output. If data errors are encountered, copies those sectors as zeroes.

**ri** as above, but if errors are encountered, attempts to recover bad ecc data with special commands.

**rh** (read headers) (SMD, IPI-2, SCSI–Fujitsu only) if the controller has a command to read sector headers, diskf reads them from tracks on the disk. If the **v** modifier is specified, diskf prints out the sector headers. On read operations, the bad sectors are returned as zeros. For example, the following command reads the sector headers for cylinder 101, head 3:

```
/etc/diskf /dev/rxdv01z rhv 101.3
```

**NOTE:** This output is vendor-dependent.

**s** (slip sector) slips a sector that has gone bad. The specified sector is slipped if the track has a spare sector. This command implies the **p** modifier.

On SCSI disks, the drive recognizes "grown" flaws as it is running, and reports the block number(s) in a message to the console. When this occurs, specify the bad block with the s command (diskf s <blocknumber>). The utility executes a "reassign block" command, which fixes the flaw and records it in the GLIST.

**S** (status) prints the label of the specified device. This label is the one read by the operating system when the drive is brought online. See the example below.

**Sa** (status-absolute) prints the label of the specified device showing the "absolute" numbers stored in the partition tables. See the example below.

**Sv** produces a disktab entry, suitable for use with dlabel. See the example below.

**t** (test arguments) prints the range of blocks specified by *from* and *to*.

**T** prints the range of blocks specified by *from* and *to* with both rounded down.

**v** (verify) verifies on a track-by-track basis that the ecc of the specified sectors is correct and reports any errors. This command takes longer than the **V** command but allows other disk requests to be serviced.

**w** (write) writes data from standard input onto the disk. Quits on write errors.

**wi** as above, but does not quit on write errors.

**x** (exercise) exercises the disk by reading and writing worst-case patterns. Full cylinder reads and writes occur if the *from* and *to* arguments are on cylinder boundaries.

**xl** as above, but does a more exhaustive test.

**z** (zeros) writes zeros to the disk.

**>** seeks from *from* to *to*. You must escape (\>) or quote (">") this command.

**?** prints on-line help. You must escape (\?) or quote ("?") this command.

### MODIFIERS

**c** (continuous mode) the given diskf command is executed until an interrupt (CTRL C) signal is received, at which time the total number of errors that the command encountered is printed. Continuous mode suppresses some informatory messages so that disk driver error messages remain on the screen. Future versions of diskf will log the errors directly from the drivers.

**e** (expert) suppresses conversion to the whole partition (z). Suppresses read/write of the internal flaw map (on SCSI drives this suppression is the default). This modifier is useful when formatting an SMD or IPI-2 disk for the first time.

**h** reads or writes headers but not data; use only with the r and w commands.

**i** ignores disk errors.

**l** specifies the long form of a test.

**n** suppresses range checking of *from* and *to*.

**p** (preserve mode) diskf tries to preserve existing data. If a sector cannot be read normally, a special command is used to read the data, even if the ecc is bad, and a diagnostic warning about corrupted data is issued.

**s** (severe mode) defeats all automatic retries and ecc correction.

**v** (verbose mode) displays more than the usual amount of information, depending on the command. For example, the f command prints the number of each cylinder formatted, rather than just the last cylinder.

### EXAMPLES: SMD AND IPI-2 DISKS

Typical disk initialization. The following commands (1) save the diskf maintained internal flaw map in a flaw file, (2) format the entire disk, slipping bad sectors indicated in the flaw file, (3) copy the flaw file to the internal flaw map on the disk, and (4) compare the actual slipped sectors on the disk with the flaw file. Formatting and comparing each take several minutes.

```
# diskf /dev/rxdv01z pe >fj1000s_99999
# diskf /dev/rxdv01z fe fj1000s_99999
# diskf /dev/rxdv01z D  fj1000s_99999
# diskf /dev/rxdv01z C internal
```

Typical disk reinitialization. The following commands (1) set the environment variable to the disk to be formatted (2) save the diskf maintained internal flaw map in a flaw file, (3) format the entire disk, slipping bad sectors indicated in the flaw file, (4) copy the flaw file to the internal flaw map on the disk, and (5) compare the actual slipped sectors on the disk with the flaw file. Formatting and comparing each take several minutes.

```
# setenv DISKF_SCRATCH_DISK /dev/rxdv01z
# diskf B >fj1000s_99999
# diskf f fj1000s_99999
# diskf D fj1000s_99999
# diskf C  internal
```

Typical track reformat, preserving data:

```
# diskf /dev/rxdv01z fp cyl.head internal
```

Typical bad sector slip:

```
# diskf /dev/rxdv01z s cyl.head.sector
```

Typical initialization of alternate track (internal is default if unspecified):

```
# diskf /dev/rxdv01z aZ
```

Sample S command output with an SMD drive:

```
# diskf /dev/rxdv00w S
magic=0x611 volume='881391635' drive='fj_m2344'
modified: Thu May 19 12:06:35 1988

        Dunix root partition: 'r' 17, swapping 'p'
Concentrix root partition: 'a' 0,
This device (/dev/rxdv00w) maj/min: 96/22 (0x60/0x16)

vdtype=0  bsize=512  cyls=624  heads=27  sectors=64  sec/cyl=1728 spare_sects=4
rpm=3600  hd/tr=0  log bits/hd=0  by/tr=40960  phys sec/trk=0 phys bsize (as)=600
cyloff=0  nblocks=1076544

L[0]=-1           L[1]=-1           L[2]=-1           L[3]=-1
L[4]=-1           L[5]=-1           L[6]=-1           L[7]=-1

P root/swap       O        C        E
b                 0c (     1c ,     0 )
d               623c (     1c ,   623 )
e                 1c (   332c ,   332 )
f               333c (   290c ,   622 )
h                 1c (   622c ,   622 )
p   Ds           44c (   127c ,   170 )
r Dr              1c (    43c ,    43 )
s                44c (   289c ,   332 )
u               171c (   135c ,   305 )
v               306c (   317c ,   622 )
w*                0c (   623c ,   622 )
x                44c (   579c ,   622 )
z                 0c (   624c ,   623 )
```

## EXAMPLES: SCSI DISKS

Typical disk initialization:

```
# dlabel -i /dev/rsd20z :bs=4k :sc=354 :nh=59 :ns=6
# diskf /dev/rsd20z f
# dlabel -iw /dev/rsd20z :bs=4k :sc=354 :nh=59 :ns=6
```

Typical set of commands to reformat a SCSI disk at a different block size after the drive has been in use for awhile:

```
diskf /dev/rsd20z p >glist     ; save current GLIST
dlabel ...                     ; change sector size
diskf /dev/rsd20z f glist      ; reformat with supplied GLIST
```

To see the GLIST which should be supplied during the format process, use the diskf p command:

```
# diskf /dev/rsd20z p
using device '/dev/rsd20z'

0        0        7347             0
0        2        22045            0
```

Sample Sa command output on a SCSI disk device:

```
Concentrix(44)# diskf /dev/rsd04z Sa
magic=0x611 volume='910572037' drive='sample'
modified: Wed Feb 27 15:54:37 1991

This device (/dev/rsd20z) maj/min: 166/119 (0xa6/0x77)
vdtype=-1  bsize=4096  cyls=1658  heads=15  sectors=12  sec/cyl=177 spare_sects=3
rpm=3600  hd/tr=1  log bits/hd=-1  by/tr=-1  phys sec/trk=0  phys bsize (as)=4096
cyloff=0  nblocks=293112

L[0]=-1          L[1]=-1          L[2]=-1          L[3]=-1
L[4]=-1          L[5]=-1          L[6]=-1          L[7]=-1

              dk_cyloff dk_nblocks dk_blkoff
b                     0       177         0
d                  1656       177    293112
f                   477    190806     84429
h                     1    292935       177
p                  1555     15222    275235
q                  1641      2655    290457
r                     1     10974       177
s                    63     73278     11151
w                     0    293112         0
z                     0    293112         0
```

Sample Sv command output on a SCSI disk device:

```
Concentrix(45)# diskf /dev/rsd04z Sv
sample:\
        :ds=sample:\
        :bs=4096:ht=1:\
        :nc=1658:nh=15:ns=12:ss=3:as=4096:rm=3600::sc=177:\
    :Ob=       0m:   :Ub=     708k: (    0.00m-    0.69m,     0.00c-    1.00c) \
    :Od=1172448k:   :Ud=1173156k: (1144.97m-1145.66m, 1656.00c-1657.00c) \
    :Of= 337716k:   :Uf=1100940k: ( 329.80m-1075.14m,  477.00c-1555.00c) \
    :Oh=     708k:  :Uh=1172448k: (    0.69m-1144.97m,     1.00c-1656.00c) \
    :Op=1100940k:   :Up=1161828k: (1075.14m-1134.60m, 1555.00c-1641.00c) \
    :Oq=1161828k:   :Uq=1172448k: (1134.60m-1144.97m, 1641.00c-1656.00c) \
    :Or=     708k:  :Ur=   44604k: (    0.69m-   43.56m,     1.00c-   63.00c) \
    :Os=   44604k:  :Us= 337716k: (   43.56m- 329.80m,    63.00c- 477.00c) \
    :Ow=       0m:  :Uw=1172448k: (    0.00m-1144.97m,     0.00c-1656.00c) \
    :Oz=       0m:  :Uz=1172448k: (    0.00m-1144.97m,     0.00c-1656.00c)
```

## NOTES

- On SCSI disks, the following command gets the "flaws needed for formatting" (the GLIST) to supply to the diskf f command:

  ```
  # diskf /dev/rsd00w p > flaw-file
  ```

  On IPI-2 disks, the disk geometry corrected version of the vendor flaws is supplied.

- On SMD disks, alternate track mapping is used when the number of flaws on any particular track exceeds the number of spare sectors. The data sectors can no longer be slipped over the bad spots and the entire track must be remapped to another track on the disk.

  Tracks mapped out are mapped to the t partition. When mapping out tracks, diskf complains and quits if the t partition does not exist, or if it is greater than two cylinders in size (implying that the partition was not set up for alternate track mapping). All FX/2800 1GB drives shipped from Alliant have a t partition set up for alternate track mapping.

## SEE ALSO

dlabel
Chapter 3, "Devices"

# dlabel

## generate disk labels from prototype file

Generates and optionally installs disk labels from a data base that resides in the file /etc/disktab. A disk label is a data structure that is stored on the first block of a disk, and contains information pertaining to the characteristics of the disk drive and the division of the physical disk drive into logical partitions.

---

**/etc/dlabel –i|–w|–o|–ob** [*option*] *label* [**/dev/r***disk***w**] [*:field=value*]
**/etc/dlabel –f** *disktab-filename*

---

### ARGUMENTS

**–b** writes the booter (/etc/booter) to disk. This is useful if you also write the label. If you only want to write the booter, use the diskf –i command.

**–Cr***r* specifies which partition (a through z), if any, is to be the root device. Do not specify more than one root device per system.

**–Cs***s*... specifies a list of all potential swap partitions for *disk*.

**–Dr***p* specifies which partition (a through z), if any, is to be the root device for Diagnostix-2800.

**–Ds***q*... specifies a list of all potential swap partitions for Diagnostix-2800.

**–f** Specifies an alternate /etc/disktab.

**–i** installs the label in the operating system. Use when the drive is not yet formatted.

**–o** outputs the disktab entry in text on standard output with all expressions resolved (this is useful for checking).

**–ob** outputs the disk label (in binary) on standard output.

**–Sr***r* same as C*r*.

**–Ss***s*... same as C*s*.

**–w** installs the label on the disk.

**–V***n* specifes an optional 16-character volume name. If omitted, dlabel generates a volume name based on the date, time, and system name.

*:field=value*

overrides the specified field of the disktab entry used. For example, :bs=4k sets the block size to 4KB.

NOTE: The following field values *must* be supplied if you have the Fujitsu M2266 SCSI disk drive: :bs=4k :ns=6 :nh=59 :sc=354

*label* specifies the prototype entry in the file /etc/disktab from which the disk label is taken.

*disk* specifies the (raw) device name of the (w partition of the) disk.

## NOTES

- If you do not specify one of the following: -i, -w, -o, or -ob, the label is not written. Previously, if you left out -i and -w, it was written to standard output.
- To relabel a disk that participates in a disk stripe, first disassemble the disk stripe with /etc/unmkstripe. After executing /etc/dlabel, remake the disk stripe with /etc/mkstripe.
- **Warning:** Do not run dlabel on active drives, not even just to print out a resolved label. Doing so may cause dlabel to change the block size on the fly; the result can be disastrous.
- Concentrix-2800 version 2.1 + supports SCSI disks. The /etc/disktab file has been altered to allow this enhancement; see that file for more information.

## EXAMPLES

The following example installs a label for a Fujitsu M2344K disk with no root partition and no paging/swapping area:

```
# /etc/dlabel -iw fj690s /dev/rxdv01z
```

The following example installs a label for a Fujitsu M2333K disk that is the root disk and holds the paging/swapping area:

```
# /etc/dlabel -iw -Crr -Css -Drp -Dsq fj690s /dev/rxdv00z
```

The root disk can be one of the following: /dev/rxdv00, /dev/rxdv01, /dev/rxdv10, or /dev/rxdv11.

The following example installs a label for a Fujitsu M2266 SCSI disk that is the root disk and holds the paging/swapping area:

```
# /etc/dlabel -iw -Crr -Css -Drp -Dsq gen_scsi /dev/rsd00z :bs=4k :ns=6 :nh=59
:sc=354
```

The following example installs a label for a Fujitsu M2266 SCSI disk with no root partition and no paging/swapping area:

```
# /etc/dlabel -iw gen_scsi /dev/rsd00z :bs=4k :ns=6 :nh=59 :sc=354
```

## FILES

/etc/disktab disk description file

## SEE ALSO

diskf
*Concentrix-2800 System Reference: Miscellaneous Routines, File Formats, and Special Files*: disktab
*Concentrix-2800 System Administation*, Chapter 3

# dmesg

## collect system diagnostic messages in error log (old)

Looks in a system buffer for recently printed diagnostic messages and prints them on the standard output. The messages are those printed by the system when device (hardware) errors occur and (occasionally) when system tables overflow non-fatally.

---

**/etc/dmesg [–] [–c]**

---

### ARGUMENTS

– computes (incrementally) the new messages since the last time *dmesg* was run and places these on the standard output.

–c prints messages continuously as they are emitted.

### FILES

/usr/adm/msgbuf scratch file for memory of – option

# dump

## incremental file system backup

Backs up a file system, totally or selectively, to magnetic tape, a disk file, or standard output. Maintains a record of the backups taken on each file system.

---

**/etc/dump** [*options* [*arguments*]] *filesystem*

---

### ARGUMENTS

The first argument is a list of options. *Options* must be specified as an unbroken character string containing keywords from the list below and defaults to **9u**. Where options take arguments (for example, **b** *factor*), the arguments follow the keyword string in corresponding order.

**0–9** specifies the dump level. Dump backs up all files that have changed since the most recent dump at a lower level. For example, if a level 2 dump is taken on Monday, a level 4 dump is taken on Tuesday, and a level 3 dump is taken on Wednesday, Wednesday's dump backs up all files that have changed since Monday.

A level 0 dump backs up the entire file system.

If no date is associated with a lower dump level, dump assumes that all files have changed.

**b** *factor* specifies the blocking factor for tape writes in even multiples of 1024 bytes. The blocking factor determines the number of bytes transferred per I/O request. A large blocking factor gives greater throughput, especially with cartridges, since larger transfers mean less time spent starting and stopping the tape drive and more time spent actually transferring data. The default blocking factor is 10 blocks (10240-byte transfers). The largest blocking factor available with some 6250 bpi tape drives is 126. A blocking factor of 32 is recommended for Exabyte cartridge tapes.

**NOTE:** Write the blocking factor on the tape label so that the same factor can be specified when restoring files.

**c** specifies that the dump medium is a cartridge instead of the standard half-inch reel. This option sets the density to 1000bpi and the length to 1700 feet. Specify the –s option (see below).

**NOTE:** Do not use the c option for Exabyte drives.

**d** *density* Specifies the density of the dump tape in bits per inch. This number is used in calculating the amount of tape used per reel. The default is 1600. (This option does not operate on Exabyte drives; density is fixed by the drive's hardware.)

**e** calculates tape or cartridge usage without actually dumping the files.

f *file* writes the dump to the specified device or file instead of the default /dev/rmt8. If the name of the file is –, dump writes to standard output.

Note: /dev/rmt8 is a link to /dev/rschs.

n notifies, by means similar to a wall command, all users in the group operator when dump requires operator attention.

Warning: This option currently does not work correctly and may cause the dump to hang. It is recommended that you not use this option until further notice.

s *size* specifies the size of the dump tape in feet. The default is 2300 feet. Dump notifies and waits for the operator to change reels based on the size and density parameters.

u writes the date of the beginning of the dump to /etc/dumpdates, if the dump completes successfully. This file records separate dates for each file system and each dump level. The format of /etc/dumpdates is plain ascii; the file contains one line for each dump of a file system (with u specified). Each line contains the name of the file system, the dump level, and the dump date.

W prints, for each file system in /etc/dumpdates, the name of the file system, its mount point, the level of the most recent dump, and the date of the most recent dump; highlights those file systems that should be dumped. All other options are ignored and dump exits immediately. The information is taken from the files /etc/dumpdates and /etc/fstab.

w is like W, but prints only those file systems that need to be dumped.

*filesystem* specifies the file system to back up.

## OPERATOR INTERVENTION

Dump logs the progress of the dump and provides other information on error output. The information includes conservative estimates of the number of blocks to write, the number of tape volumes required, the time to completion, and the time to a tape volume change. Requests for operator interventation are made to error output and replies are expected on standard input. If the n option is specified, the requests for operator intervention are also sent to all users in the operator group.

Dump requires operator intervention on these conditions: end of tape, end of dump, tape write error, tape open error, or disk read error (above the threshold of 32). All responses to dump are y (yes) or n (no).

Checkpoints are taken at the start of each tape volume. If writing that volume fails for some reason, you can rewind the bad tape volume, mount a new tape volume, and restart the dump at the beginning of the volume.

## SUGGESTED SCHEDULE

Shut the system down to single-user mode while taking a dump.

Take level 0 dumps of each file system periodically.

Take incremental dumps of active file systems daily. Use a set of at least 10 tapes per file system on a cyclical basis. For small file systems, daily level 1 dumps may be practical. For larger file systems, take level 1 dumps weekly using a set of at least five tapes per file system on a cyclical basis and using the following "modified Tower of Hanoi" sequence for daily dumps:

```
3 2 5 4 7 6
```

After each level 1 dump, the sequence repeats, beginning with level 3.

## EXAMPLES

The following command performs a level 0 dump of the file system /dev/rxdv00f. If the dump completes successfully, /etc/dumpdates is updated. The blocking factor for Exabyte tape writes is 32 (corresponds to b). The tape density is 6250 bpi (6250 corresponds to d). The size of the tape is specified as 30,000 blocks (corresponds to s).The output device is the Exabyte drive rschs (corresponds to f).

```
# /etc/dump 0ubdsf 32 6250 30000 /dev/rschs /dev/rxdv00f
```

The following command dumps the file system /dev/rxdv00r to the file /tmp/dump_root.

```
# /etc/dump 0ubdsf 32 6250 30000 /dev/rschs /dev/rxdv00r
```

## ERRORS

A return status of 0 indicates success. A return status of 1 indicates a start-up error. A return status of 3 indicates abnormal termination.

## NOTES

- Fewer than 32 read errors on the file system are ignored.
- Each volume requires a new process; parent processes for volumes already written sleep until the entire tape is written.
- The w and W options do not report file systems never before recorded in /etc/dumpdates even if listed in /etc/fstab.

## FILES

| | |
|---|---|
| /dev/rmt8 | default tape unit to dump to |
| /etc/dumpdates | new format dump date record |
| /etc/fstab | dump table: file systems and frequency |
| /etc/group | to find group operator |

## SEE ALSO

rdump, restore, rrestore
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: dump, fstab
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: ctime

# dumpfs

## dump file system information

Prints the superblock and cylinder group information for a specified file system. The listing is long and detailed.

---

## /etc/dumpfs *filesys*

---

### ARGUMENTS

*filesys* name of the file system. The name can be the mount point, the raw device name, or the block device name.

### DESCRIPTION (4.2/4.3BSD FILE SYSTEM)

The superblock information is as follows:

| | |
|---|---|
| 0x011954 | magic number |
| time | last time written |
| sblkno | address of superblock |
| cblkno | offset of cylinder block |
| iblkno | offset of i-node blocks |
| dblkno | offset of first data after cylinder group |
| sbsize | actual size of superblock |
| cgsize | cylinder group size |
| cgoffset | cylinder group offset in cylinder |
| cgmask | used to calculate mod of fs_ntrak |
| ncg | number of cylinder groups |
| size | number of fragments in file system |
| blocks | number of data fragments in file system |
| bsize | size of basic blocks in file system |
| bshift | blkno calculation of logical block number |
| bmask | blkoff calculation of block offsets |
| fsize | size of fragmentation blocks in file system |
| fshift | numfrags calculation of number of fragments |
| fmask | fragoff calculation of fragment offsets |
| frag | number of fragments in a block |
| fragshift | block to fragment shift |
| fsbtodb | fsbtodb and dbtofsb shift constant |
| minfree | minimum percentage of free blocks |
| maxbpg | maximum number of blocks per cylinder group |
| maxcontig | maximum number of contiguous blocks |
| rotdelay | number of milliseconds for optimal next block |
| rps | disk revolutions per second |
| csaddr | block address of cylinder group summary area |
| cssize | size of cylinder group summary area |
| csshift | csum block number |
| csmask | csum block offset |

| | |
|---|---|
| ntrak | tracks per cylinder |
| nsect | sectors per track |
| spc | sectors per cylinder |
| ncyl | cylinders in file system |
| cpg | cylinders per group |
| bpg | blocks per group |
| fpg | fragments per group |
| ipg | i-nodes per group |
| nindir | value of NINDIR |
| inopb | value of INOPB |
| nspf | value of NSPF |
| id | file system identification |
| nbfree | number of free blocks |
| ndir | number of directories |
| nifree | number of free i-nodes |
| nffree | number of free frags |
| cgrotor | last cylinder group searched |
| fmod | superblock modified flag |
| ronly | mounted read-only flag |

The cylinder group information is as follows:

| | |
|---|---|
| 0x090255 | magic number |
| time | time last written |
| cgx | cylinder group number (offset) |
| ncyl | number of cylinders in this cylinder group |
| niblk | number of i-node blocks in this cylinder group |
| ndblk | number of data blocks in this cylinder group |
| nbfree | number of free blocks |
| ndir | number of directories |
| nifree | number of free i-nodes |
| nffree | number of free frags |
| rotor | position of last used block |
| irotor | position of last used i-node |
| frotor | position of last used fragment |
| frsum | counts of available fragments |
| iused | used i-node map |
| free | free block map |
| b | positions of free blocks |

## DESCRIPTION (TRACKED FILE SYSTEM)

The superblock information is as follows:

| | |
|---|---|
| 0x081486 | magic number |
| time | last time written |
| sblkno | address of superblock |
| cblkno | offset of cylinder block |
| iblkno | offset of i-node blocks |
| dblkno | offset of first data after cylinder group |
| sbsize | actual size of superblock |
| cgsize | cylinder group size |
| cgoffset | cylinder group offset in cylinder |
| cgmask | used to calculate mod of fs_ntrak |
| ncg | number of cylinder groups |
| size | number of blocks in file system |
| blocks | number of data blocks in file system |
| nstripes | number of partitions in file system (1 means unstriped) |
| bsize | size of basic blocks in file system |
| fsize | size of fragmentation blocks in file system |
| fshift | numfrags calculation of number of fragments |
| fmask | fragoff calculation of fragment offsets |
| frag | number of fragments in a block |
| fsbtodb | fsbtodb and dbtofsb shift constant |
| fbmmask | fragment bit map: block mask right justified |
| fullblock | first block that is not fragmentable |
| minfree | minimum percentage of free blocks |
| reserved | number of reserved fragments |
| maxbpg | maximum number of blocks per cylinder group |
| csaddr | block address of cylinder group summary area |
| cssize | size of cylinder group summary area |
| indirsize | indirect block size |
| ntrak | tracks per cylinder |
| nsect | sectors per track |
| spc | sectors per cylinder |
| ncyl | cylinders in file system |
| cpg | cylinders per group |
| bpg | blocks per group |
| fpg | fragments per group |
| ipg | i-nodes per group |
| nindir | value of NINDIR |
| id | file system identification |
| nbfree | number of free blocks |
| ndir | number of directories |
| nifree | number of free i-nodes |
| nffree | number of free frags |
| fmod | superblock modified flag |
| ronly | mounted read–only flag |
| fbmshift | fragment bit map shift factors |
| fbmshmsk | fragment bit map in place block masks |
| stripesz | fragments in each partition |
| firstcg | first cylinder group in each partition |
| cgrotor | last cylinder group searched |

The cylinder group information is as follows:

| | |
|---|---|
| 0x082586 | magic number |
| time | time last written |
| cgx | cylinder group number (offset) |
| stripe | stripe number (offset) |
| base | base disk address of this cylinder group |
| ndblk | number of data blocks in this cylinder group |
| niblk | number of i-node blocks in this cylinder group |
| nbfree | number of free blocks |
| ndir | number of directories |
| nifree | number of free i-nodes |
| rotor | position of last used block |
| irotor | position of last used i-node |
| frotor | position of last used fragment |
| frsum | counts of available fragments |
| iused | used i-node map |
| free | free block map |
| b | positions of free blocks |

## SEE ALSO

fsck, makefs, tunefs
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: fs, disktab

# dumpterm

print contents of compiled terminfo file in human-readable form

Dumpterm reads the given files and decodes their contents to derive a reasonable representation of the terminfo entry which produced the file originally. It should be noted, in case of emergency, that the output of dumpterm is perfectly usable as the input to compile.

---

**/usr/lib/terminfo/dumpterm** *file* ...

---

## SEE ALSO

compile
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: term, terminfo

# edquota

## edit user quotas

Edits quotas for selected users, or edits the soft time limits. For each user, or for the soft time limits, a temporary file is created with an ASCII representation of the current disk quotas or time limits, and an editor is then invoked on the file. The quotas may then be modified. After you exit the editor, edquota reads the temporary file and modifies the binary quota files to reflect the changes made.

The editor invoked is vi unless the environment variable EDITOR specifies otherwise.

---

**/etc/edquota** [**-p** *proto-user*] *usernames ...*
**/etc/edquota** [**-t**]

---

### ARGUMENTS

**-p** duplicate the quotas of the prototypical user and do not invoke the editor. This is the normal mechanism used to initialize quotas for groups of users.

**-t** edit the soft time limits for each file system. If the time limits are zero, the default time limits in <ufs/quota.h> are used. Time units of sec(onds), min(utes), hour(s), day(s), week(s), and month(es) are understood. Time limits are printed in the greatest possible time unit such that the value is greater than or equal to one.

### DESCRIPTION

Only the superuser can edit quotas. For quotas to be established on a file system, the root directory of the file system must contain a file, owned by root, called quotas. See quotaon for details.

For each *username* that you edit, the temporary quota file contains a line of the following form for each file system containing a quotas file:

```
fs /usr blocks (soft =       0K, hard =       0K) inodes (soft = 0, hard = 0)
```

A specification of 0 (the default) imposes no limits. Adjust the soft and hard limits for blocks and inodes to suit the user. Adjust only the numeric specifications and leave the rest of the text as is. Each numeric specification must be an integer optionally followed by the letter K (kilobytes), M (megabytes), or G (gigabytes). The blocks limit specifies the amount of space the user can use in the file system; the inodes limit specifies the number of files the user can have in the file system. The soft limit causes a warning when it is exceeded (within a time limit; see the time limit discussion below); the hard limit stops any operation that would exceed the limit.

The following example permits any number of files (inodes) but limits the user to 10MB of space in /usr. A warning is generated if the user exceeds 8MB.

```
fs /usr blocks (soft =       8M, hard =      10M) inodes (soft = 0, hard = 0)
```

When the soft limit is exceeded, the user has a certain amount of time to reduce usage on the file system below the soft limit. If the soft limit remains in excess over the time limit, the user is excluded from further activity on the file system until the usage is brought within the soft limit. The default time limit specified in <ufs/quota.h> as of Concentrix-2800 2.2 is seven days. You can adjust this limit by invoking edquota –t. The temporary file is of the following format:

```
fs /usr blocks time limit = 0 (default), files time limit = 0 (default)
```

Change the numeric specification and substitue for "(default)" a time unit. The following example gives the user two days to reduce usage on /usr below the soft limit after it has been exceeded:

```
fs /usr blocks time limit = 2.00 days, files time limit = 2.00 days
```

## FILES

| | |
|---|---|
| quotas | quota file at the file system root |
| /etc/mtab | mounted file systems |

## SEE ALSO

quotacheck, quotaon, quotaoff, repquota, rquotad
*Concentrix-2800 Commands and Applications Manual*: quota, vi
*Concentrix-2800 System Reference: System Calls and C Library Routines:* quotactl, rquota

# exportfs

## export and unexport directories to NFS clients

Makes a local directory (or file) available for mounting over the network by NFS clients. Typically invoked at boot time by the /etc/rc.local script using information contained in the /etc/exports file. The superuser can run exportfs at any time to alter the list or the characteristics of exported directories. Directories that are currently exported are listed in the file /etc/xtab.

---

**/usr/etc/exportfs [-aiuv] [-o *options*] [*directory*]**

---

### ARGUMENTS

Running the command with no arguments lists the currently exported directories.

**-a** (all) exports all directories listed in /etc/exports; or, if -u is specified, unexports all currently exported directories.

**-i** ignore the options in /etc/exports. By default, exportfs consults /etc/exports for the options associated with the exported directory.

**-u** unexports the indicated directories. By default, directories are exported.

**-v** (verbose) prints the name of each directory as it is exported or unexported.

**-o** specifies a comma-separated list of optional characteristics for the directory being exported:

**ro** exports the directory read-only. By default, the directory is exported read-write.

**rw=***hostname*[:*hostname*]...
exports the directory read-write only to the specified hosts, read-only to other machines.

**anon=***uid* sets *uid* as the effective user ID for requests from unknown (anonymous) users. Super-users (uid 0) are considered unknown unless included in the root option below. Defaults to -2. Setting anon to -1 disables anonymous access. Note that by default secure NFS accepts insecure requests as anonymous, and those wishing for extra security can disable this feature by setting "anon" to -1.

**root=***hostname*[:*hostname*]...
gives root access only to the root users from the specified hosts. Defaults to no root access.

**access=***client*[:*client*]...
specifies the clients that can mount a directory. *Client* can be either a host name or a net group. A name in /etc/netgroup takes precedence over a name in /etc/hosts. By default, any client can mount a directory.

**secure** Require clients to use a more secure protocol when accessing the directory.

*directory* name of directory or file to be exported; must be specified as a full pathname.

## FILES

| | |
|---|---|
| /etc/exports | static export information |
| /etc/netgroup | list of network groups |
| /etc/xtab | current state of exported directories |

## NOTES

You cannot export a directory that is either a parent or a sub-directory of a directory that is currently exported and within the same filesystem. For example, you cannot export both /usr and /usr/local if both directories reside in the same disk partition.

## SEE ALSO

*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: exports, netgroup

# extracting

## extract from a tar formatted software distribution tape

Extracts a particular software release file from a distribution tape. Intended to be called from setup_exec as part of the installation procedure for diskless/NFS client support. The release file is assumed to be in tar format and is read from the tape into the current working directory.

---

**/usr/etc/install/extracting** *tapedev skip blockfactor keywords* [*tapeserver*]

---

### ARGUMENTS

These arguments must be specified in the order given in the synopsis above.

*tapedev* is the full pathname of the non-rewinding raw tape device on which the distribution tape is mounted. If *tapedev* is on a remote machine, the optional *tapeserver* argument is the host name of the remote machine.

*skip* is the number of tape files to skip over before starting to read (this argument is passed to the command mt to position the tape).

*blockfactor* is the blocking factor to use with tar while reading the file.

*keywords* is a quoted string that is used as part of the output to the user to identify what is being loaded from tape.

### SEE ALSO

INSTALL, setup_exec, verify_tapevol_arch
*Concentrix-2800 Commands and Applications Manual*: mt, rsh, tar

# extendswap

## enable/disable multiple-disk swapping/paging area

Enables or disables a multi-partition paging and swapping area.

---

**/usr/alliant/extendswap** [*flag*] [*kernel-name*]

---

### ARGUMENTS

*flag* one of two values:

**off** (or OFF) disable a multi-partition paging and swapping area.

**on** (or ON; default) enable a multi-partition paging and swapping area.

*kernel-name*
(defaults to /unix) name of the kernel image.

### DESCRIPTION

By default, Concentrix-2800 searches for and selects a single swapping/paging partition at boot time. Extendswap causes Concentrix-2800 to select all partitions that are marked as potential swapping partitions (see dlabel) as part of the swapping/paging area.

### NOTES

- The set of partitions to be used for a multiple-disk swapping/paging area must be on disks and controllers all of the same type.
- At least one swapping/paging partition must be accessible to the master IP (IP0).
- Although swapping and paging are single-threaded (the system does not use multiple controllers concurrently as in disk I/O), multiple controllers provide some throughput enhancement.

### EXAMPLES

The following example enables a multi-partition paging and swapping area:

```
# /usr/alliant/extendswap on
```

### SEE ALSO

dlabel

# fastboot

## reboot system without checking disks

Reboots the system, without checking the file systems, or performs other actions depending on the arguments. Logs the reboot unless -l, -n, or -q is specified.

---

**/etc/fastboot** [*argument* ...]

---

### ARGUMENTS

**-B** stops in VMEbus booter mode and prompts for names for the Diagnostix-2800 kernel, the Diagnostix-2800 kernel debugger, and the Diagnostix-2800 root file system. Type RETURN to continue with default file names. You can skip loading the Diagnostix-2800 debugger by typing ESCAPE. From the booter prompt, you can also type %h to get help and %ls to list the root directory. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-b** stops in FX/2800 booter mode and prompts for file names for the Concentrix-2800 kernel, the Concentrix-2800 kernel debugger, and the processor configuration file. Type RETURN to continue with default file names. From the booter prompt, you can also type ! (exclamation mark) to connect to Diagnostix-2800 (type CTRL d to return to the FX/2800 booter). This option causes only Concentrix-2800 to be rebooted.

**-c** same as -H.

**-D** stops in Alliant ddt mode and waits for commands. You can proceed by typing ESCAPE p or ESCAPE g. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-d** stops in Alliant ddt 2800 mode and waits for commands. You can proceed by typing ESCAPE p or ESCAPE g. This option causes only Concentrix-2800 to be rebooted.

**-H** halts the system (shuts the system down to console mode). Recommended for power-down. From the console, type u to exit from console mode and start the system. From a remote terminal, type o r RETURN u.

**-h** shuts the system down to Diagnostix-2800 mode. To start Concentrix-2800, type CTRL d, then respond y to the FX/2800 prompt.

**-l** inhibits the log message.

**-n** does not sync the disks (for emergencies only); inhibits the log message.

**-q** reboots quickly, without first shutting down running processes; inhibits the log message.

**-S** stops in Diagnostix-2800 mode. To continue the reboot, type CTRL d, then respond y to the FX/2800 prompt. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

-s stops at single-user mode in Concentrix-2800. To continue the reboot, type sync[RETURN], /etc/umount -a[RETURN], sync[RETURN], and [CTRL]d. This option causes only Concentrix-2800 to be rebooted.

## DESCRIPTION

Fastboot is a shell script that creates a file /fastboot and invokes the reboot command. The existence of the file /fastboot causes the file system checking to be skipped.

## EXAMPLES

The following example reboots the system without checking the file systems.

```
# /etc/fastboot
```

The following example reboots the system and stops in Alliant ddt mode.

```
# /etc/fastboot -D
```

The following example reboots the system and stops in Diagnostix-2800 mode.

```
# /etc/fastboot -S
```

## FILES

| | |
|---|---|
| /dunix | Diagnostix kernel image |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /vmunix | a symbolic link to /unix |

## SEE ALSO

fsck, init, rc, reboot, shutdown

# fasthalt

## halt system without checking disks

Shuts the system down to Diagnostix-2800 mode (no -H argument) or console mode (-H argument) and causes file system checking to be skipped on the next reboot. Logs the shutdown in the system log unless -l, -n, or -q is specified.

---

**/etc/fasthalt [-H] [-l] [-n] [-q] [-y]**

---

### ARGUMENTS

**-H** halts the system (stops in console mode). Recommended for power-down. From the console, type u to exit from console mode and resume system startup. From a remote terminal, type o r[RETURN]u.

**-l** inhibits the log message.

**-n** does not sync the disks (for emergencies only); inhibits the log message.

**-q** reboots quickly, without first shutting down running processes; inhibits the log message.

**-y** halts the system from a dial-up port.

### DESCRIPTION

If -H is not specified, Concentrix-2800 shuts down and the system is placed in Diagnostix-2800 mode. To restart Concentrix-2800, type [CTRL]d, then respond y to the FX/2800 prompt.

You do not have to halt the system to work in Diagnostix-2800 mode. You can escape from Concentrix-2800 on the console by typing [CTRL]^. (control-shift-6 followed by a period). Then respond X to the FX/2800 startup prompt. To return to Concentrix-2800, type [CTRL]d, then respond t to the FX/2800 startup prompt.

Fasthalt is a shell script that creates a file /fastboot and invokes the halt command. The existence of the file /fastboot causes file system checking to be skipped.

### FILES

/fastboot empty file whose existence signals a fastboot

# fingerd

## remote user information server

Provides an interface to the finger program.

The fingerd daemon is started as needed by /etc/inetd.

---

## /etc/fingerd

---

### DESCRIPTION

Fingerd is a simple protocol based on RFC742. The program returns a friendly, human-oriented status report on either the system at the moment or on a particular person in depth. No format is required and the protocol consists mostly of specifying a single command line.

Fingerd listens for TCP requests at port 79. Once connected, fingerd reads a single command line terminated by a newline, which is passed to the finger command. Fingerd closes its connections as soon as the output is finished.

If the line is null (that is, just a newline), fingerd returns a default report that lists all people logged into the system at that moment.

If a user name is specified, the response lists more extended information for only that particular user, whether logged in or not. Allowable names in the command line include both login names and usernames. If a name is ambiguous, all possible derivations are returned.

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: finger

# fix_bootparams

## add or delete entries from /etc/bootparams

Adds or removes a diskless/NFS client's entries from the server's /etc/bootparams file. This file specifies the NFS pathname of the client's root partition, and swap and dump files. Fix_bootparams is intended to be run by setup_client as part of the administration of diskless/NFS clients.

---

**/usr/etc/install/fix_bootparams** *op clientname* [*rootpath swappath dumppath*]

---

### OPTIONS

The pathname arguments are only required for the add operation and are specified in standard NFS format: *servername:pathname_on_server*.

*op* is the operation to be performed:

**add** add entries.

**remove** remove entries.

*clientname* is the hostname of the client to be added or removed.

*rootpath* is the NFS pathname of the client's root partition.

*swappath* is the NFS pathname of the client's swap file.

*dumppath* is the NFS pathname of the client's crash dump file. May be specified as none.

### FILES

/etc/bootparams

### SEE ALSO

INSTALL, setup_client
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: bootparams

# fsck

## file system consistency check / interactive repair

Performs consistency checks on and repairs file systems. In non-interactive mode, makes minor repairs automatically. In interactive mode, makes repairs based on your responses. A file system must not be active during its repair; best practice is to always execute /etc/fsck in single-user mode.

---

**/etc/fsck -p** [*filesys* ...]

**/etc/fsck** [*option* ...] [*filesys* ...]

---

### ARGUMENTS

**-b** *block* uses the specified block as the superblock for the file system. Block 24 or 32 (old file systems) is always an alternate superblock.

**-B** *list* displays information about the specified list of fragment numbers. The format of *list* is one or more numbers or ranges of numbers (hyphenated pairs), separated by commas. (See "Logical Block Number" below.) The -B option displays the following information for each specified fragment number. In phase one, from left to right: the physical fragment number (within the partition), the number of fragments in the block, the logical block number within the i-node (@ indicates an indirect pointer block; @@ indicates a doubly-indirect pointer block), and the i-node numbers. In phase two, from left to right: i-node numbers and the pathnames that refer to each i-node.

**-I** *list* displays information about the specified list of i-nodes. The format of a list is one or more numbers or ranges of numbers (hyphenated pairs), separated by commas. The -I option displays the following information for each specified i-node. In phase one, from left to right: the physical fragment number (within the partition), the number of fragments in the block, the logical block number within the i-node (@ indicates an indirect pointer block; @@ indicates a doubly indirect pointer block), and the i-node number itself. In phase two, from left to right: the i-node number and the pathnames that refer to that i-node.

**-n** assumes a no response to all questions asked by fsck; does not open the file system for writing.

**-p** invokes the non-interactive mode of fsck (see below): The -p option cannot be used with any of the other options.

**-y** assumes a yes response to all questions asked by fsck; use this option with caution as it allows fsck to continue no matter what problems are encountered.

*filesys* name of the file system. The name can be the mount point, the raw device name, or the block device name. Defaults to all the file systems in /etc/fstab.

## LOGICAL BLOCK NUMBER

The file system disk block number (fsb) displayed by a disk error message refers to 512-byte disk blocks. These must be converted to a fragment as follows for use with the -B option of fsck (standard 4.2BSD has 1024-byte fragments; tracked file sytems have 4096-byte fragments):

- Standard 4.2BSD file system - divide by 2.
- Tracked file system - divide by 8.

## NON-INTERACTIVE MODE

In non-interactive (-p) mode, /etc/fsck performs a consistency check on the specified file systems. Typically, the system initialization script /etc/rc executes /etc/fsck -p on all file systems during an automatic reboot, unless a fastboot is in effect. Disks are processed in parallel to check the file systems as quickly as possible.

The file /etc/fstab specifies on which pass to check each file system. Typically, the root file system is checked on pass 1, other root file systems are checked on pass 2, other small file systems are checked on separate passes, and finally the large user file systems are checked on the last pass. Only partitions that have a non-zero pass number and that are mounted rw or ro are checked.

In non-interactive (-p) mode, /etc/fsck corrects only the following inconsistencies:

- Unreferenced i-nodes
- Link counts in i-nodes too large
- Missing blocks in the free list
- Blocks in the free list also in files
- Wrong counts in the superblock

If more serious inconsistencies are discovered, /etc/fsck exits with an abnormal return status (an automatic reboot fails) and /etc/fsck must be executed in interactive mode. For each corrected inconsistency, one or more lines are printed identifying the file system and the nature of the correction. After successfully correcting a file system, /etc/fsck prints the number of files on that file system, the number of used and free blocks, and the percentage of fragmentation.

If you send a QUIT signal (kill command) to fsck, the file system checks complete but fsck exits with an error status. This action is useful during a reboot to complete the file system checks, then stop in single-user mode.

## INTERACTIVE MODE

In interactive mode, /etc/fsck prompts for concurrence before a correction is attempted. Corrective action for conditions not fixable under /etc/fsck -p may result in loss of data. Examine the diagnostic output to determine the amount and severity of the data lost. The default action for each consistency correction is to wait for you to respond y or n. If you do not have write permission, the default action is -n.

Inconsistencies checked are as follows:

1. Blocks claimed by more than one i-node or the free list.
2. Blocks claimed by an i-node or the free list outside the range of the file system.
3. Incorrect link counts.
4. Size checks: directory size not of proper format.
5. Bad i-node format.
6. Blocks not accounted for anywhere.
7. Directory checks: file pointing to unallocated i-node; i-node number out of range.
8. Superblock checks: more blocks for i-nodes than exist in the file system.
9. Bad free block list format.
10. Total free block and/or free i-node count incorrect.

Orphaned files and directories (allocated but unreferenced) are, with your concurrence, reconnected by placing them in the lost+found directory. The name assigned is the i-node number. If the directory lost+found does not exist, it is created; if its size is insufficient, its size is increased.

## DESCRIPTION

A quiescent file system is checked for structural integrity by performing consistency checks on the redundant data intrinsic to the file system. The redundant data is either read from the file system or computed from other known values. The file system *must* be in a quiescent state.

### Superblock Checking

The most commonly corrupted item in a file system is the summary information associated with the superblock. The summary information is prone to corruption because it is modified with every change to the file system's blocks or i-nodes, and is usually corrupted after an unclean halt.

The superblock is checked for inconsistencies involving file system size, number of i-nodes, free block count, and free i-node count. The file system size must be larger than the number of blocks used by the superblock and the number of blocks used by the list of i-nodes. The file system size and layout information are the most critical pieces of information for /etc/fsck. While there is no way to actually check the sizes, since they are statically determined by /etc/makefs, /etc/fsck can check that these sizes are within reasonable bounds. All other file system checks require that these sizes be correct. If /etc/fsck detects corruption in the static parameters of the default superblock, /etc/fsck requests the operator to specify the location of an alternate superblock.

### Free Block Checking

All blocks marked as free in the cylinder group block maps must not be claimed by files. When all blocks are initially accounted for, /etc/fsck checks that the number of free blocks plus the number of blocks claimed by the i-nodes equals the total number of blocks in the file system.

If anything is wrong with the block allocation maps, /etc/fsck rebuilds them based on the list of allocated blocks.

The summary information associated with the superblock counts the total number of free blocks within the file system. This count is compared to the number of free blocks within the file system. If the two counts do not agree, the incorrect count in the summary information is replaced by the actual free block count.

The summary information counts the total number of free i-nodes within the file system. Fsck compares this count to the number of free i-nodes it found within the file system. If the two counts do not agree, the incorrect count in the summary information is replaced by the actual free i-node count.

### Checking the I-node State

An individual i-node is not as likely to be corrupted as the allocation information. However, because of the great number of active i-nodes, a few of the i-nodes are usually corrupted.

The list of i-nodes in the file system is checked sequentially starting with i-node 2 (i-node 0 marks unused i-nodes; i-node 1 is reserved) and progressing through the last i-node in the file system. The state of each i-node is checked for inconsistencies involving format and type, link count, duplicate blocks, bad blocks, and i-node size.

Each i-node contains a mode word. This mode word describes the type and state of the i-node. i-nodes must be one of six types: regular, directory, symbolic link, special block, special character, or socket. i-nodes may be in one of three allocation states: unallocated, allocated, and neither. The last state suggests an incorrectly formatted i-node; an i-node can get in this state if bad data is written into the i-node list. The only possible corrective action is to clear the i-node.

### I-node Links

Each i-node counts the total number of directory entries linked to the i-node.The link count of each i-node is calculated starting at the root of the file system and descending through the directory structure.

If the stored link count is non-zero and the actual link count is zero (no directory entry appears for the i-node), the disconnected file is placed in the lost+found directory. If the stored and actual link counts are non-zero and unequal, a directory entry may have been added or removed without the i-node being updated; the incorrect stored link count is replaced by the actual link count.

Each i-node contains a list, or pointers to lists (indirect blocks), of all the blocks claimed by the i-node. Since indirect blocks are owned by an i-node, inconsistencies in indirect blocks directly affect the i-node that owns it.

Each block number claimed by an i-node is compared to a list of already allocated blocks. If another i-node already claims a block number, the block number is added to a list of duplicate blocks. Otherwise, the list of allocated blocks is updated to include the block number.

If duplicate blocks exist, /etc/fsck performs a partial second pass over the i-node list to find the i-node of the duplicated block. If this condition does arise, the i-node with the earliest modify time is usually incorrect and should be cleared. You are prompted and must decide which i-node to keep and which to clear.

The range of each block number claimed by an i-node is checked. If the block number is lower than the first data block in the file system, or greater than the last data block, the block number is a bad block number. Many bad blocks in an i-node are usually caused by an indirect block that was not written to the file system. If an i-node contains bad block numbers, you are prompted to clear it.

### I-node Data Size

Each i-node contains a count of the number of its data blocks. The number of actual data blocks is the sum of the allocated data blocks and the indirect blocks. The actual number of data blocks are compared to the actual number of blocks the i-node claims. If an i-node contains an incorrect count, you are prompted to fix the count.

Each i-node records the number of data bytes in the file associated with the i-node. The consistency of the byte size field is roughly checked by computing from the size field the maximum number of blocks that should be associated with the i-node, and comparing that expected block count against the actual number of blocks the i-node claims.

### Checking the Data Associated with an I-node

An i-node can directly or indirectly reference three kinds of data blocks. Plain data blocks contain the information stored in a file; symbolic link data blocks contain the path name stored in a link; directory data blocks contain directory entries. All referenced blocks must be the same kind. Only the validity of directory data blocks are checked.

Each directory data block is checked for several types of inconsistencies. These inconsistencies include directory i-node numbers pointing to unallocated i-nodes, directory i-node numbers that are greater than the number of i-nodes in the file system, incorrect directory i-node numbers for . (current directory) and .. (parent directory), and directories that are not attached to the file system. If the i-node number in a directory data block references an unallocated i-node, the directory entry is removed.

If a directory entry i-node number references outside the i-node list, the directory entry is removed. This condition occurs if bad data is written into a directory data block.

The directory i-node number entry for . (current directory) must be the first entry in the directory data block and must reference itself, that is, must equal the i-node number for the directory data block. The directory i-node number entry for .. (parent directory) must be the second entry in the directory data block and must equal the i-node number for the parent of the directory entry (or the i-node number of the directory data block if the directory is the root directory). Incorrect directory i-node numbers are replaced with the correct values.

### File System Connectivity

The general connectivity of the file system is checked. If directories are not linked into the file system, the directory is linked back into the file system in the lost+found directory.

## ERROR CONDITIONS

The error conditions are organized by the phase in which they can occur. The error conditions that may occur in more than one phase are discussed in initialization.

### Initialization

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 initialization errors are fatal in non-interactive mode.

*option* option?
: A specified option is not legal; legal options are -b, -B, -I, -n, -p, and -y. Fsck terminates.

cannot alloc *n* bytes for blockmap
: A request for virtual memory failed. Call Alliant Customer Service. Fsck terminates.

cannot alloc *n* bytes for freemap
: A request for virtual memory failed. Call Alliant Customer Service. Fsck terminates.

cannot alloc *n* bytes for statemap
: A request for virtual memory failed. Call Alliant Customer Service. Fsck terminates.

cannot alloc *n* bytes for lncntp
: A request for virtual memory failed. Call Alliant Customer Service. Fsck terminates.

Can't open checklist file: *F*
: The file system checklist file F (usually /etc/fstab) can not be opened for reading. Fsck terminates on this error condition. Check access modes of F.

Can't stat root
: Fsck's request for statistics about the root directory "/" failed. This should never happen. Fsck terminates on this error condition. Call Alliant.

Can't stat *filesys*
: A request for statistics about the file system *filesys* failed. In interactive mode, the file system is skipped. Check the access modes of *filesys*.

Can't make sense out of name *filesys*
: A request for statistics about the file system *filesys* failed. In interactive mode, the file system is skipped. Check the access modes of *filesys*.

Can't open *filesys*
: *An* attempt to open the file system *filesys* failed. In interactive mode, the file system is skipped. Check the access modes of *filesys*.

*filesys*: (NO WRITE)
Either the -n flag was specified or an attempt to open the file system *filesys* for writing failed. In interactive mode, all diagnostics are printed but no modifications are attempted.

file is not a block or character device; OK?
You specified a regular file by mistake. Check the type of the file specified.

YES Ignore this error condition.

NO ignore this file system and continues checking the next file system given.

*filesys*: BAD SUPER BLOCK: B
The superblock has been corrupted. An alternative superblock must be selected from among those listed when the file system was created. For file systems with a blocksize less than 32K, specifying -b 24 is a good first choice. This message is usually preceded by one of the following messages:

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

and is followed by the message:

USE -b OPTION TO FSCK TO SPECIFY LOCATION OF AN ALTERNATE SUPER-BLOCK TO SUPPLY NEEDED INFORMATION; SEE fsck(8).

INTERNAL INCONSISTENCY: *message*
An internal panic occurred. Call Alliant Customer Service.

CAN NOT SEEK: BLK *n* (CONTINUE)?
Moving to the specified block in the file system failed. You should probably call Alliant Customer Service.

YES attempt to continue to run the file system check. Often, however the problem persists. This error condition does not allow a complete check of the file system. A second run of /etc/fsck should be made to recheck this file system. If the block was part of the virtual memory buffer cache, /etc/fsck terminates with the message, "Fatal I/O error."

NO terminate the program.

CAN NOT READ: BLK *b* (CONTINUE)?
Reading the specified block in the file system failed. You should probably call Alliant Customer Service.

YES attempt to continue to run the file system check. Often, however, the problem persists. This error condition does not allow a complete check of the file system. A second run of /etc/fsck should be made to recheck this file system. If the block was part of the virtual memory buffer cache, /etc/fsck terminates with the message, "Fatal I/O error."

NO terminate the program.

CAN NOT WRITE: BLK *n* (CONTINUE)?
Writing the specified block number in the file system failed. The disk is write-protected. You should probably call Alliant Customer Service.

YES attempt to continue to run the file system check. Often, however, the problem persists. This error condition does not allow a complete check of the file system. A second run of /etc/fsck should be made to re-check this file system. If the block was part of the virtual memory buffer cache, /etc/fsck terminates with the message, "Fatal I/O error."

NO terminate the program.

### Phase 1 – Check Blocks and Sizes

This section lists error conditions resulting from checking i-node types, setting up the zero-link-count table, examining i-node block numbers for bad and duplicate blocks, checking i-node size, and checking i-node format. All errors in this phase, except INCORRECT BLOCK COUNT, are fatal in non-interactive mode.

CG *cg*: BAD MAGIC NUMBER
The magic number of cylinder group *cg* is wrong. This usually indicates that the cylinder group maps have been destroyed. In interactive mode, the cylinder group is marked as needing to be reconstructed.

UNKNOWN FILE TYPE I=*i-node* (CLEAR)?
The mode word of the specified i-node indicates that the i-node is not a special block i-node, special character i-node, socket i-node, regular i-node, symbolic link, or directory i-node.

YES deallocate i-node *i-node* by zeroing its contents. This action always invokes the UNALLOCATED error condition in Phase 2 for each directory entry pointing to this i-node.

NO ignore this error condition.

LINK COUNT TABLE OVERFLOW (CONTINUE)?
An internal table containing allocated i-nodes with a link count of zero has no more room. Report the error to Alliant Customer Service.

YES continue with the program. This error condition does not allow a complete check of the file system. A second run should be made to recheck this file system. If another allocated i-node with a zero link count is found, this error condition is repeated.

NO terminate the program.

*block* BAD I=*i-node*
I-node *i-node* contains block number *block* with a number lower than the number of the first data block in the file system or greater than the number of the last block in the file system. This error condition may invoke the EXCESSIVE BAD BLKS error condition in Phase 1 if i-node *i-node* has too many block numbers outside the file system range. This error condition always invokes the BAD/DUP error condition in Phase 2 and Phase 4.

EXCESSIVE BAD BLKS I=*i-node* (CONTINUE)?
I-node *i-node* has more than a tolerable number (usually 10) of blocks with a number lower than the number of the first data block in the file system or greater than the number of last block in the file system associated with i-node I.

YES ignore the rest of the blocks in this i-node and continue checking with the next i-node in the file system. This error condition does not allow a complete check of the file system. A second run should be made to recheck this file system.

NO terminate the program.

*block* DUP I=*i-node*
I-node *i-node* contains block number *block*, which is already claimed by another i-node. This error condition may invoke the EXCESSIVE DUP BLKS error condition in Phase 1 if i-node *i-node* has too many block numbers claimed by other i-nodes. This error condition always invokes Phase 1b and the BAD/DUP error condition in Phase 2 and Phase 4.

EXCESSIVE DUP BLKS I=*i-node* (CONTINUE)?
More than a tolerable number (usually 10) of blocks are claimed by other i-nodes.

YES ignore the rest of the blocks in this i-node and continue checking with the next i-node in the file system. This error condition does not allow a complete check of the file system. A second run should be made to recheck this file system.

NO terminate the program.

DUP TABLE OVERFLOW (CONTINUE)?
An internal table in fsck containing duplicate block numbers has no more room. Report this error to Alliant Customer Service.

YES continue with the program. This error condition does not allow a complete check of the file system. A second run should be made to recheck this file system. If another duplicate block is found, this error condition repeats.

NO terminate the program.

PARTIALLY ALLOCATED INODE I=*i-node* (CLEAR)?
I-node *i-node* is neither allocated nor unallocated.

YES deallocate i-node I by zeroing its contents.

NO ignore this error condition.

INCORRECT BLOCK COUNT I=*i-node* (*x* should be *y*) (CORRECT)?
The block count for i-node *i-node* is *x* blocks, but should be *y* blocks. In non-interactive mode, the count is corrected.

YES replace the block count of i-node *i-node* with *y*.

NO ignore this error condition.

### Phase 1B: Rescan for More Dups

When a duplicate block is found in the file system, the file system is rescanned to find the i-node that previously claimed that block.

*block* DUP I=*i-node*
I-node *i-node* contains block number *block*, which is already claimed by another i-node. This error condition always invokes the BAD/DUP error condition in Phase 2. You can determine which i-nodes have overlapping blocks by examining this error condition and the DUP error condition in Phase 1.

### Phase 2 – Check Pathnames

This section lists error conditions resulting from root i-node mode and status, directory i-node pointers in range, and directory entries pointing to bad i-nodes. All errors in this phase are fatal in non-interactive mode.

ROOT INODE UNALLOCATED. TERMINATING.
The root i-node (usually i-node number 2) has no allocate mode bits. This should never happen. The program terminates.

NAME TOO LONG *filename*
A pathname is excessively long. This condition is usually indicative of loops in the file system name space and can occur if the super-user has made circular links to directories. The offending links must be removed.

ROOT INODE NOT DIRECTORY (FIX)?
The root i-node (usually i-node number 2) is not directory i-node type.

YES replace the root i-node's type to be a directory. If the root i-node's data blocks are not directory blocks, a VERY large number of error conditions occur.

NO terminate the program.

DUPS/BAD IN ROOT INODE (CONTINUE)?
Phase 1 or Phase 1b found duplicate blocks or bad blocks in the root i-node (usually i-node number 2) for the file system.

YES ignore the error condition in the root i-node and attempt to continue to run the file system check. If the root i-node is not correct, a large number of other error conditions may occur.

NO terminate the program.

I OUT OF RANGE I=*i-node* NAME=*filename* (REMOVE)?
A directory entry has an i-node number that is greater than the end of the i-node list.

YES remove the directory entry.

NO ignore this error condition.

UNALLOCATED
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (REMOVE)?
A directory entry has a directory i-node without allocate mode bits. The owner, mode, size, modify time, and directory name are printed.

YES remove the directory entry.

NO ignore this error condition.

DUP/BAD I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (REMOVE)?
Phase 1 or Phase 1b found duplicate blocks or bad blocks associated with directory entry *f*, directory i-node *i-node*. The owner, mode, size, modify time, and directory name are printed.

YES remove the directory entry.

NO ignore this error condition.

ZERO LENGTH DIRECTORY
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (REMOVE)?
A directory entry has a size that is zero. The owner, mode, size, modify time, and directory name are printed.

YES remove the directory entry; this action always invokes the BAD/DUP error condition in Phase 4.

NO ignore this error condition.

DIRECTORY TOO SHORT
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory's size is less than the minimum size directory. The owner, mode, size, modify time, and directory name are printed.

YES increase the size of the directory to the minimum directory size.

NO ignore this directory.

DIRECTORY CORRUPTED
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (SALVAGE)?
A directory has an inconsistent internal state.

YES throw away all entries up to the next 512-byte boundary. This rather drastic action can throw away up to 42 entries, and should be taken only after other recovery efforts have failed.

NO Skip up to the next 512-byte boundary and resume reading, but do not modify the directory.

BAD INODE NUMBER FOR '.'
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory's i-node number for . (current directory) does does not equal *i-node*.

YES change the i-node number to *i-node*.

NO leave the i-node number unchanged.

MISSING '.' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory's first entry is unallocated.

YES make an entry for . (current directory) with i-node number equal to *i-node*.

NO leave the directory unchanged.

MISSING '.' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f*
CANNOT FIX, FIRST ENTRY IN DIRECTORY CONTAINS *f*
A directory's first entry is *f*. Fsck cannot resolve this problem. The file system should be mounted and the offending entry moved elsewhere. The file system should then be unmounted and fsck should be run again.

MISSING '.' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f*
CANNOT FIX, INSUFFICIENT SPACE TO ADD '.'
A directory's first entry is not . (current directory). Fsck cannot resolve this problem as it should never happen. Call Alliant Customer Service.

EXTRA '.' ENTRY I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory has more than one entry for . (current directory).
YES remove the extra entry.

NO leave the directory unchanged.

BAD INODE NUMBER FOR '..'
I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory's i-node number for .. (parent directory) does not equal the parent of *i-node*.

YES change the i-node number to the parent of *i-node*.

NO leave the i-node number unchanged.

MISSING '..' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory's second entry is unallocated.

YES make an entry for .. (parent directory) with i-node number equal to the parent of *i-node*.

NO leave the directory unchanged.

MISSING '..' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f*
CANNOT FIX, SECOND ENTRY IN DIRECTORY CONTAINS *f*
A directory's second entry is *f*. Fsck cannot resolve this problem. The file system should be mounted and the offending entry moved elsewhere. The file system should then be unmounted and fsck should be run again.

MISSING '..' I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f*
CANNOT FIX, INSUFFICIENT SPACE TO ADD '..'
A directory's second entry is not .. (parent directory). Fsck cannot resolve this problem as it should never happen. Call Alliant Customer Service.

EXTRA '..' ENTRY I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* DIR=*f* (FIX)?
A directory has more than one entry for .. (parent directory).

YES remove the extra entry.

NO leave the directory unchanged.

### Phase 3 - Check Connectivity

This section lists error conditions resulting from unreferenced directories, and missing or full lost+found directories.

UNREF DIR I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (RECONNECT)?
The directory i-node *i-node* was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of directory i-node *i-node* are printed. In non-interactive mode, the directory is reconnected if its size is non-zero; otherwise it is cleared.

YES reconnect directory i-node *i-node* to the file system in the directory for lost files (usually lost+found). This action may invoke the lost+found error condition in Phase 3 if problems occur connecting directory i-node *i-node* to lost+found. This action may also invoke the CONNECTED error condition in Phase 3 if the link was successful.

NO ignore this error condition. This action always invokes the UNREF error condition in Phase 4.

SORRY. NO lost+found DIRECTORY
No lost+found directory is in the root directory of the file system; fsck ignores the request to link a directory in lost+found. This error always invokes the UNREF error condition in Phase 4. Check the access modes of lost+found if it exists. This error is fatal in non-interactive mode.

SORRY. NO SPACE IN lost+found DIRECTORY
No space exists 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 error always invokes the UNREF error condition in Phase 4. Clean out unnecessary entries in lost+found or make lost+found larger. This error is fatal in non-interactive mode.

DIR I=*i-node1* CONNECTED. PARENT WAS I=*i-node2*
This is an advisory message indicating that directory i-node *i-node1* was successfully connected to the lost+found directory. The parent i-node *i-node2* of the directory i-node *i-node1* is replaced by the i-node number of the lost+found directory.

### Phase 4 – Check Reference Counts

This section lists error conditions resulting from unreferenced files; missing or full lost+found directories; incorrect link counts for files, directories, symbolic links, and special files; unreferenced files, symbolic links, and directories; bad and duplicate blocks in files, symbolic links, and directories; and incorrect total free i-node counts. All errors in this phase are correctable in non-interactive mode except running out of space in the lost+found directory.

UNREF FILE I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (RECONNECT)?
i-node *i-node* was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of i-node *i-node* are printed. In non-interactive mode, the file is cleared if either its size or its link count is zero; otherwise it is reconnected.

YES reconnect i-node *i-node* to the file system in the directory for lost files (usually lost+found). This action may invoke the lost+found error condition in Phase 4 if problems occur connecting i-node *i-node* to lost+found.

NO ignore this error condition. This action always invokes the CLEAR error condition in Phase 4.

UNREF FILE I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (CLEAR)?
The i-node mentioned in the immediately previous error condition cannot be reconnected. This condition cannot occur in non-interactive mode, since lack of space to reconnect files is a fatal error.

YES deallocate the i-node mentioned in the immediately previous error condition by zeroing its contents.

NO ignore this error condition.

SORRY. NO lost+found DIRECTORY
No lost+found directory exists in the root directory of the file system; fsck ignores the request to link a file in lost+found. This error always invokes the CLEAR error condition in Phase 4. Check the access modes of lost+found. This error is fatal in non-interactive mode.

SORRY. NO SPACE IN lost+found DIRECTORY
No space exists 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 error always invokes the CLEAR error condition in Phase 4. Check the size and contents of lost+found. This error is fatal in non-interactive mode.

LINK COUNT FILE I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* COUNT=*x* SHOULD BE y (ADJUST)?
The link count for i-node *i-node*, which is a file, is *x* but should be *y*. The owner, mode, size, and modify time are printed. In non-interactive mode, the link count is adjusted.

YES replace the link count of file i-node *i-node* with *y*.

NO ignore this error condition.

LINK COUNT DIR I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* COUNT=*x* SHOULD BE *y* (ADJUST)?
The link count for i-node *i-node*, which is a directory, is *x* but should be *y*. The owner, mode, size, and modify time of directory i-node *i-node* are printed. In non-interactive mode, the link count is adjusted.

YES replace the link count of directory i-node *i-node* with *y*.

NO ignore this error condition.

LINK COUNT F I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* COUNT=*x* SHOULD BE *y* (ADJUST)?
The link count for F i-node *i-node* is *x* but should be *y*. The name, owner, mode, size, and modify time are printed. In non-interactive mode, the link count is adjusted.

YES replace the link count of i-node *i-node* with *y*.

NO ignore this error condition.

UNREF FILE I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (CLEAR)?
I-node *i-node*, which is a file, was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of i-node *i-node* are printed. In non-interactive mode, the i-node is cleared since it is a file that was not connected because its size or link count was zero.

YES deallocate i-node *i-node* by zeroing its contents.

NO ignore this error condition.

UNREF DIR I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (CLEAR)?
I-node *i-node*, which is a directory, was not connected to a directory entry when the file system was traversed. The owner, mode, size, and modify time of i-node *i-node* are printed. In non-interactive mode, the i-node is cleared since it is a directory that was not connected because its size or link count was zero.

YES deallocate i-node *i-node* by zeroing its contents. Run fsck again.

NO ignore this error condition.

BAD/DUP FILE I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (CLEAR)?
Phase 1 or Phase 1b found duplicate blocks or bad blocks associated with file i-node *i-node*. The owner, mode, size, and modify time of i-node *i-node* are printed. This error cannot arise in non-interactive mode, as it would have caused a fatal error earlier.

YES deallocate i-node *i-node* by zeroing its contents.

NO ignore this error condition.

BAD/DUP DIR I=*i-node* OWNER=*o* MODE=*m* SIZE=*s* MTIME=*t* (CLEAR)?
Phase 1 or Phase 1b found duplicate blocks or bad blocks associated with directory i-node *i-node*. The owner, mode, size, and modify time of i-node *i-node* are printed. This error cannot arise in non-interactive mode, as it would have caused a fatal error earlier.

YES deallocate i-node *i-node* by zeroing its contents.

NO ignore this error condition.

FREE INODE COUNT WRONG IN SUPERBLK (FIX)?
The actual count of the free i-nodes does not match the count in the superblock of the file system. In non-interactive mode, the count is fixed.

YES replace the count in the superblock by the actual count.

NO ignore this error condition.

### Phase 5 – Check Cylinder Groups

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 being incorrect.

CG *cg*: BAD MAGIC NUMBER
The magic number of cylinder group *cg* is wrong. This error usually indicates that the cylinder group maps have been destroyed. In interactive mode, the cylinder group is marked as needing to be reconstructed. This error is fatal in non-interactive mode.

EXCESSIVE BAD BLKS IN BIT MAPS (CONTINUE)?
An i-node contains more than a tolerable number (usually 10) of blocks claimed by other i-nodes or that are out of the legal range for the file system. This error is fatal in non-interactive mode.

YES ignore the rest of the free block maps and continue execution.

NO terminate the program.

SUMMARY INFORMATION *t* BAD
Where *t* is one or more of:

(INODE FREE)
(BLOCK OFFSETS)
(FRAG SUMMARIES)
(SUPER BLOCK SUMMARIES)

The indicated summary information is incorrect. This error always invokes the BAD CYLINDER GROUPS error in Phase 6. In non-interactive mode, the summary information is recomputed.

*n* BLK(S) MISSING
The specified number of blocks, unused by the file system, were not in the free block maps. This error always invokes the BAD CYLINDER GROUPS error in Phase 6. In non-interactive mode, the block maps are rebuilt.

BAD CYLINDER GROUPS (SALVAGE)?
Phase 5 found bad blocks in the free block maps, duplicate blocks in the free block maps, or blocks missing from the file system. In non-interactive mode, the cylinder groups are reconstructed.

YES replace the actual free block maps with a new free block maps.

NO ignore this error condition.

FREE BLK COUNT WRONG IN SUPERBLOCK (FIX)?
The actual count of free blocks does not match the count in the superblock of the file system. In non-interactive mode, the counts are fixed.

YES replace the count in the superblock by the actual count.

NO ignore this error condition.

### Phase 6 – Salvage Cylinder Groups

This phase reconstructs the free block maps. No error messages are produced.

### Phase 7 – Check Stripes

This phase checks the stripe allocation summaries.

STRIPE SUMMARIES BAD (FIX)?
Phase 7 found a bad stripe summary. In non-interactive mode, the stripe summary is reconstructed.

YES replace the actual stripe summary with the new stripe summary.

NO ignore this error condition.

ALLOCATION SUMMARY BAD (FIX)?
Phase 7 found the allocation summary to be incorrect. In non-interactive mode, the allocation summary is reconstructed.

YES replace the allocation summary with the new allocation summary.

NO ignore this error condition.

### Cleanup

This section lists advisory messages about the file system and the modify status of the file system.

*v* files, *w* used, *x* free (*y* frags, *z* blocks)
This is an advisory message indicating that the file system contains *v* files using *w* fragment size blocks leaving *x* fragment size blocks free in the file system. The numbers in parentheses breaks the free count down into *y* free fragments and *z* free full sized blocks.

***** REBOOT UNIX *****
This is an advisory message indicating that the root file system has been modified. If the system is not rebooted immediately, the work done by fsck may be undone by the in-core copies of tables the operating system keeps. In non-interactive mode, fsck exits with a code of 4. The auto-reboot script interprets an exit code of 4 by issuing a reboot system call.

***** FILE SYSTEM WAS MODIFIED *****
This is an advisory message indicating that the current file system was modified. If this file system is mounted or is the current root file system, fsck should be halted and the system rebooted. If the system is not rebooted immediately, the work done by fsck may be undone by the in-core copies of tables the system keeps.

## EXAMPLE

```
# /etc/fsck -I4233 -n /dev/rxdv10h
** /dev/rxdv10h (NO WRITE)
** Last Mounted on /rev2
** Phase 1 - Check Blocks and Sizes
   28671   1       0   i4233
** Phase 2 - Check Pathnames
inode   4233   /lib/crt0.o
inode   4233   /lib/libcrt0.a
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cylinder groups
** Phase 7 - Check Stripes
2542 files, 28761 used, 7394 free (6 frags, 1847 blocks)
```

```
/etc/fsck -B28668-28671 -n /dev/rxdv10h
** /dev/rxd10p (NO WRITE)
** Last Mounted on /rev2
** Phase 1 - Check Blocks and Sizes
   28668   3       0   i4213
   28671   1       0   i4233
** Phase 2 - Check Pathnames
inode   4213   /usr/include/dfif/init.h
inode   4233   /lib/crt0.o
inode   4233   /lib/libcrt0.a
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Cylinder groups
** Phase 7 - Check Stripes
2542 files, 28761 used, 7394 free (6 frags, 1847 blocks)
```

## FILES

| | |
|---|---|
| /etc/fstab | contains default list of file systems to check |

## SEE ALSO

makefs, mklost+found, reboot
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: fstab, fs

# fsirand

## install random i-node generation numbers

The fsirand command helps increase the security of filesystems exported by NFS. It installs random i-node generation numbers on all the i-nodes on a specified device and installs a filesystem ID in the superblock.

---

**/etc/fsirand [-p]** *special*

---

### ARGUMENTS

**-p** prints out but does not change the generation numbers for all the i-nodes.

*special* specifies a device.

### DESCRIPTION

Use fsirand only on an unmounted filesystem that has been checked with fsck. It can, however, be used on the root filesystem in single-user mode, if the system is immediately rebooted afterwords.

### SEE ALSO

fsck

# ftpd

## DARPA Internet file transfer protocol server

Uses the TCP protocol and listens at the port specified in the ftp service as specified in /etc/services.

The ftpd daemon is started as needed by /etc/inetd.

---

## /etc/ftpd [-d] [-l] [-t*timeout*]

---

### ARGUMENTS

-d writes debugging information to the system log.

-l logs each ftp session in the system log.

-t*timeout* sets the inactivity timeout period to *timeout* (default is 15 minutes).

### DESCRIPTION

The ftp server currently supports the following ftp requests; case is not distinguished.

| Request | Description |
|---|---|
| ABOR | abort previous command |
| ACCT | specify account (ignored) |
| ALLO | allocate storage (vacuously) |
| APPE | append to a file |
| CDUP | change to parent of current working directory |
| CWD | change working directory |
| DELE | delete a file |
| HELP | give help information |
| LIST | give list files in a directory (`ls -lg`) |
| MKD | make a directory |
| MODE | specify data transfer *mode* |
| NLST | give name list of files in directory (`ls`) |
| NOOP | do nothing |
| PASS | specify password |
| PASV | prepare for server-to-server transfer |
| PORT | specify data connection port |
| PWD | print the current working directory |
| QUIT | terminate session |
| RETR | retrieve a file |
| RMD | remove a directory |
| RNFR | specify rename-from file name |
| RNTO | specify rename-to file name |
| STOR | store a file |
| STOU | store a file with a unique name |

| | |
|---|---|
| STRU | specify data transfer *structure* |
| TYPE | specify data transfer *type* |
| USER | specify user name |
| XCUP | change to parent of current working directory |
| XCWD | change working directory |
| XMKD | make a directory |
| XPWD | print the current working directory |
| XRMD | remove a directory |

The remaining ftp requests specified in Internet RFC 959 are recognized but not implemented.

The ftp server aborts an active file transfer only when the ABOR command is preceded by a telnet IP (interrupt process) signal and a telnet synch signal in the command telenet stream, as described in Internet RFC 959.

Ftpd interprets file names according to the "globbing" conventions used by the C shell, which allows use of the metacharacters *?[]{}-.

Ftpd authenticates users according to the following rules:

1. The user name must be in the password data base /etc/passwd and not have a null password. In this case a password must be provided by the client before any file operations can be performed.
2. The user name must not appear in the file /etc/ftpusers.
3. The user must have a standard shell returned by the C routine getusershell.
4. If the user name is "anonymous" or "ftp," an anonymous ftp account must be present in the password file (user ftp). The user is allowed to log in by specifying any password (by convention, the password is the client host's name).

In the last case, ftpd takes special measures to restrict the client's access privileges. The server performs a chroot command to the home directory of the ftp user. In order that system security is not breached, it is recommended that the ftp subtree be constructed with care; the following rules are recommended:

| | |
|---|---|
| ~ftp | Make the home directory owned by ftp and unwritable by anyone. |
| ~ftp/bin | Make this directory owned by the super-user and unwritable by anyone. The program ls must be present to support the list commands. This program should have mode 111. |
| ~ftp/etc | Make this directory owned by the super-user and unwritable by anyone. The files passwd and group must be present for the ls command to work properly. These files should be mode 444. |
| ~ftp/pub | Make this directory mode 777 and owned by ftp. Users should place files that are to be accessible via the anonymous account in this directory. |
| ~ftp/lib | Make this directory owned by the superuser and unwritable by anyone. The file c.shared must be present for the ls command in ~ftp/bin to work. It can be copied from /lib/c.shared. This file should have mode 111. (This directory must contain any shared libraries used by commands accessible to the anonymous ftp user in order for those commands to successfully execute.) |

## NOTES

- The anonymous account is inherently dangerous and should be avoided when possible.

- The server must run as the super-user to create sockets with privileged port numbers. It maintains an effective user id of the logged in user, reverting to the super-user only when binding addresses to sockets. The possible security holes have been extensively scrutinized, but are possibly incomplete.

## SEE ALSO

syslogd
*Concentrix-2800 Commands and Applications Manual*: ftp
*Concentrix-2800 System Reference*: *System Calls and C Routines*: getusershell
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services, passwd, group

# gettable

## get NIC format host tables from a host

**NOTE:** This software is provided on an "as is" basis without warranty expressed or implied. Any software services available will be provided at the then current charges.

Obtains the NIC standard host tables from a "nicname" server. The specified *host* is queried for the tables. The tables, if retrieved, are placed in the file *output-file* or hosts.txt.

---

**/etc/gettable** [**-v**] *host* [*output-file*]

---

### ARGUMENTS

**-v** gets just the version number and places it in *output-file* or hosts.ver.

### DESCRIPTION

Gettable operates by opening a TCP connection to the port indicated in the service specification for nicname. A request is then made for ALL names and the resultant information is placed in the output file.

Gettable is best used in conjunction with the htable program, which converts the NIC standard file format to that used by the network library lookup routines.

### SEE ALSO

htable
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: gethostent, getnetent

# getty

## set terminal mode

Determines the characteristics of a terminal connected to the system, prompts for and obtains a login name from the terminal, and initiates a login process for the terminal.

---

**/etc/getty** [*type* [*tty*]]

---

### DESCRIPTION

Typically, init calls getty with an argument that is used as an index into the file /etc/gettytab to determine the characteristics of the line. If the argument or /etc/gettytab is missing, a default table is used.

The argument *tty* specifies the name of a file in /dev to open for the terminal. If the argument is omitted or is a hyphen (-), the tty line is opened as file descriptor 0.

Getty performs the following actions by default. Most of the default actions of getty can be modified by adjusting /etc/gettytab.

1. Prints a banner that includes the name of the system and prompts for a login name.

2. Reads the login name until a a newline or carriage return occurs. The latter results in the system being set to treat carriage returns appropriately.

3. If the login name contains no lowercase alphabetic characters (and the name is non-empty), flags the system to map all upper-case characters into the corresponding lower-case characters.

4. If a null character (assumed to be the result of the user pushing BREAK), the speed is changed and the login: prompt is typed again; a second BREAK changes the speed again and the login: is typed again. Successive BREAK characters cycle through the same standard set of speeds.

5. Invokes the login command with the login name as an argument.

Getty can be set to time out after some interval, which causes dial-up lines to hang up if the login name is not entered reasonably quickly.

## DIAGNOSTICS

- `ttyxx: No such device or address`
- `ttyxx: No such file or address`

  A terminal that is turned on in the file /etc/ttys cannot be opened. The probable reasons: the tty line is not configured into the system; the device was not attached at boot time; the special file in /dev does not exist.

## FILES

| | |
|---|---|
| /etc/gettytab | terminal configuration data base |
| /etc/ttys | terminal initialization data |

## SEE ALSO

init
*Concentrix-2800 Commands and Applications Manual*: login
*Concentrix-2800 System Reference: System Calls and C Library Routines*: ioctl
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: gettytab, ttys

# halt

## halt system

Shuts the system down to Diagnostix-2800 mode (no -H argument) or console mode (-H argument). Logs the shutdown in the system log unless -l, -n, or -q is specified.

---

**/etc/halt [-H] [-l] [-n] [-q] [-y]**

---

### ARGUMENTS

-H halts the system (shuts the system down to console mode). Recommended for power-down. From the console, type u to exit from console mode and resume system startup. From a remote terminal, type o r[RETURN]u.

-l inhibits the log message.

-n does not sync the disks (for emergencies only); inhibits the log message.

-q reboots quickly, without first shutting down running processes; inhibits the log message.

-y halts the system from a dial-up port.

### DESCRIPTION

If -H is not specified, Concentrix-2800 shuts down and the system is placed in Diagnostix-2800 mode. To restart Concentrix-2800, type [CTRL]d, then respond y to the FX/2800 prompt.

You do not have to halt the system to work in Diagnostix-2800 mode. You can escape from Concentrix-2800 on the console by typing [CTRL]^. (control-shift-6 followed by a period). Then respond X to the FX/2800 startup prompt. To return to Concentrix-2800, type [CTRL]d, then respond t to the FX/2800 startup prompt.

# htable

## convert NIC standard format host tables

**NOTE:** This software is provided on an "as is" basis without warranty expressed or implied. Any software services available will be provided at the then current charges.

Converts host files in the format specified in Internet RFC 810 to the format used by the network library routines. Three files are created as a result of running htable: hosts, networks, and gateways.

---

**/etc/htable** [**-c** *connected-nets*] [**-l** *local-nets*] *host*

---

### DESCRIPTION

The hosts file may be used by the gethostent routines in mapping host names to addresses if the name server named is not used. The *networks* file is used by the getnetent routines in mapping network names to numbers. The gateways file may be used by the routing daemon in identifying "passive" Internet gateways; see routed for an explanation.

If any of the files localhosts, localnetworks, or localgateways are present in the current directory, the file's contents are prepended to the output file; of these, only the gateway file is interpreted. This allows sites to maintain local aliases and entries that are not normally present in the master database. Only one gateway to each network is placed in the gateways file; a gateway listed in the localgateways file overrides any in the input file.

If the gateways file is to be used, a list of networks to which the host is directly connected is specified with the option -c. The networks, separated by commas, may be given by name or in Internet-standard dot notation, for example, arpanet,128.32,local-ether-net. Htable only includes gateways that are directly connected to one of the networks specified, or that can be reached from another gateway on a connected net.

If the -l option is specified with a list of networks (in the same form as -c), these networks are treated as "local," and information about hosts on local networks is taken only from the localhosts file. Entries for local hosts from the main database are omitted. This allows the localhosts file to completely override any entries in the input file.

Htable is best used in conjunction with the gettable program that retrieves the NIC database from a host.

### SEE ALSO

gettable, named
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: gethostent, getnetent

# ifconfig

## configure network interface parameters

Assigns an address to a network interface and/or configures network interface parameters.The system start-up scripts typically invoke ifconfig to define the network address of each interface present on a machine. Ifconfig can be used after start-up to redefine an interface's address or other operating paramters.

Only the super-user can modify the configuration of a network interface.

---

**/etc/ifconfig** *interface address-family* [*address*] [*dest-address*] [*parameters*]

**/etc/ifconfig** *interface* [*protocol-family*]

---

### ARGUMENTS

Ifconfig displays the current configuration for a network interface when no optional parameters are supplied. If a protocol family is specified, ifconfig limits the display to that protocol family.

*interface* is a string of the form *name-unit*, for example: ex0, that specifies an ethernet controller.

*address-family*
is inet (Internet).

*address* is either a host name present in the host name data base /etc/hosts, or a DARPA Internet address expressed in the Internet standard "dot notation."

The host number can be omitted on 10mb/s Ethernet interfaces that use the hardware physical address, and on interfaces other than the first.

If ypbind is running, the local /etc/hosts file is not consulted. The host name must exist in the /etc/hosts file on the ypserver.

*parameters* is one or more of the following parameters:

**arp** enables the use of the Address Resolution Protocol in mapping between network level addresses and link level addresses (default). This is currently implemented for mapping between DARPA Internet addresses and 10Mb/s Ethernet addresses.

**-arp** disables the use of the Address Resolution Protocol.

**broadcast** (inet only) specifies the address to use to represent broadcasts to the network. The default broadcast address is the address with a host part of all ones.

**debug** enables driver dependent debugging code. Typically, this option turns on extra console error logging.

**-debug** disables driver dependent debugging code.

**down** marks an interface down. When an interface is marked down, the system does not attempt to transmit messages through that interface. If possible, the interface is reset to disable reception as well. This action does not automatically disable routes using the interface.

**dstaddr** specifies the address of the correspondent on the other end of a point-to-point link.

**metric** *n* sets the routing metric of the interface to *n* (default 0). The routing metric is used by the routing protocol routed. Higher metrics make a route less favorable; metrics are counted as additional hops to the destination network or host.

**netmask** *mask*
(inet only) specifies how much of the address to reserve for dividing networks into sub-networks. The mask includes the network part of the local address and the subnet part, which is taken from the host field of the address. The maks can be specified as a single hexadecimal number with a leading 0x, with a dot notation internet address, or with a pseudo-network name listed in the network table networks. The mask contains ones for the bit positions in the 32-bit address that are to be used for the network and subnet parts, and zeros for the host part. The mask should contain at least the standard network portion, and the subnet field should be contiguous with the network portion.

**trailers** enables the use of a "trailer" link level encapsulation when sending (default). If a network interface supports trailers, the system, when possible, encapsulates outgoing messages in a manner that minimizes the number of memory-to-memory copy operations performed by the receiver. On networks that support the Address Resolution Protocol (arp), trailers indicates that the system should request that other systems use trailers when sending to this host. Similarly, trailer encapsulations are sent to other hosts that have made such requests.

**-trailers** disables the use of a "trailer" link level encapsulation.

**up** marks an interface up. Enables an interface previously marked as down. Automatic when setting the first address on an interface. If the interface was reset when previously marked down, the hardware is reinitialized.

## DESCRIPTION

The system start-up script /etc/rc.local, as delivered, invokes ifconfig to configure Ethernet controller 0, unit 0.

```
ifconfig ex0 -trailers netmask my-netmask broadcast my-broadcast 'hostname'
```

You must supply appropriate values for my-netmask and my-broadcast, or delete these parameters, for example:

```
ifconfig ex0 -trailers broadcast 125.0.0.0 `hostname`
```

Only the superuser can modify the configuration of a network interface.

## DIAGNOSTICS

Messages indicating: the specified interface does not exit; the requested address is unknown; the user is not privileged; and tried to alter an interface's configuration.

## SEE ALSO

*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: hosts
*Concentrix-2800 Commands and Applications Manual*: netstat

# ifstats

## print out ethernet board statistics

Prints statistics kept by the Excelan or CMC-130 ethernet board, such as the number of buffers received and transmitted, the number of buffers lost due to lack of resources, and so forth.

---

**/usr/alliant/ifstats** *interface* [-r]

---

### ARGUMENTS

*interface* is a string of the form *name unit*, for example: en0, that specifies an ethernet controller.

-r resets the on-board statistics to 0. Requires super-user privilege.

### EXAMPLE

```
# /usr/alliant/ifstats en0
Transmit data:
frames sent without errors:                          1657725
frames sent despite SQE test errors:                 1475830
```

*The IEEE 802.3 specification says that after an Ethernet frame is transmitted, the transceiver shall assert the Signal Quality Error signal to exercise the collision detect circuitry. This is also called the "heartbeat" mechanism. This signal is not generated by older Ethernet transceivers and some repeaters. Non-zero are important if your hardware should be generating SQE. In this case, the collision detection mechanism may be defective which will lead to very poor network performance.*

```
frames sent after deferral due to active medium:      33356
```

*This notes the number of times the Ethernet controller had a frame to send but had to wait because some other system was transmitting.*

```
frames sent after a single collision:                  5946
```

*This counts the number of times this controller and some other system started to transmit a message within 64 microseconds of each other. After waiting a moment, the controller was then able to send the message without difficulty.*

```
frames sent after multiple collisions:                 3386
```

*Same as above, but 2 to 15 collisions ocurred before the frame made it out.*

```
frames abandoned after 16 collisions:                     0
```

*After 16 attempts, Ethernet gives up and discards the message.*

```
frames abandoned due to late collision:                   1
```

*A late collision is a collision detected more than 64 microseconds after a frame begins trasmission. This should not occur in a properly designed Ethernet. A network that is too large will produce these errors, as can noise on the coax cable.*

```
frames abandoned due to no carrier:                    0
```

*This usually means the the controller is not plugged into the transceiver. It may also indicate a short in the coax.*

```
frames abandoned due to length > 1518:                 0
frames abandoned due to silo underflow:                0
frames abandoned due to board level memory error:      0
```

*These three counters should remain zero.*

```
Receive data:
frames received without errors:                  1987635
frames received with CRC errors:                     753
```

*CRC errors indicates frames were corrupted. The Ethernet driver discards them.*

```
frames received with alignment errors:               669
```

*Alignment errors indicate frames were received with a length that was not a multiple of eight bits. The Ethernet driver discards them.*

```
frames lost due to no receive buffers:                77
```

*This indicates messages that arrived faster than they could be processed.*

```
frames lost due to silo overrun:                       0
frames lost due to board level memory error:           0
```

*These counts should remain 0.*

A subset of the above is printed for the ex device:

```
# ifstats ex0
Transmit data:
frames sent without errors:                            0
frames sent despite SQE test errors:                   0
frames abandoned after 16 collisions:                  0

Receive data:
frames received without errors:                        0
frames received with CRC errors:                       0
frames received with alignment errors:                 0
frames lost due to no receive buffers:                 0
```

## SEE ALSO

ifconfig

# implog

## IMP log interpreter

**NOTE:** This software is provided on an "as is" basis without warranty expressed or implied. Any software services available will be provided at the then current charges.

Implog interprets the message log produced by *implogd*.

---

**/etc/implog** [*option* ...]

---

### ARGUMENTS

If no arguments are specified, implog interprets and prints every message present in the message file. The following options can be specified to force printing only a subset of the logged messages:

**-D** suppresses data messages.

**-f** follows the logging process in action. *Implog* prints the current contents of the log file, then checks for new logged messages every five seconds.

**-c** shows the contents of the data in hexadecimal bytes, in addition to printing any data messages logged,

**-l** *link#* shows only those messages received on the specified "link". If no value is given for the link, the link number of the IP protocol is assumed.

**-h** *host#* shows only those messages received from the specified host. (Usually specified in conjunction with an imp.)

**-i** *imp#* shows only those messages received from the specified imp.

**-r** Prints the raw imp leader, showing all fields, in addition to the formatted interpretation according to type.

**-t** *message-type*
shows only those messages received of the specified message type.

### NOTES

- Implog cannot specify multiple hosts, imps, etc. and cannot follow reception of messages without looking at those currently in the file.

### SEE ALSO

implogd
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* imp

# implogd

## IMP logger process

> **NOTE:** This software is provided on an "as is" basis without warranty expressed or implied. Any software services available will be provided at the then current charges.

Logs error messages from the IMP, placing them in the file /usr/adm/implog.

---

## /etc/implogd

---

### DESCRIPTION

Entries in the file are variable length. Each log entry has a fixed length header of the form:

```
struct sockstamp {
        short  sin_family;
        u_short       sin_port;
        struct in_addr sin_addr;
        time_t        sin_time;
        int    sin_len;
};
```

followed, possibly, by the message received from the IMP. Each time the logging process is started up it places a time stamp entry in the file (a header with *sin_len* field set to 0).

The logging process catches only those message from the IMP that are not processed by a protocol module, for example, IP. This implies the log should contain only status information such as IMP going down messages, host down and other error messages, and, perhaps, stray NCP messages.

### SEE ALSO

implog
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: imp

# inetd

## internet services daemon

Invokes the internet server processes listed in /etc/inetd.conf as needed. Other servers start from direct invocation, typically in the system start-up files rc and rc.local.

---

**/etc/inetd [-d] [*configuration-file*]**

---

### ARGUMENTS

**-d** specifies that debugging traces are to be turned on for connection-oriented (TCP) services.

*configuration-file*
specifies a file containing configuration information; defaults to /etc/inetd.conf. See below for a description.

### DESCRIPTION

Connection-oriented (TCP) services are invoked each time a connection is made by creating a process. The process is passed the connection as file descriptor 0 and an argument of the form *sourcehost.sourceport* where *sourcehost* is hexadecimal and *sourceport* is decimal.

Datagram-oriented (UDP) services are invoked when a datagram arrives; a process is created and passed the connection as file descriptor 0. Inetd looks at the socket where datagrams arrive again only after this process completes. The paradigms for such processes are either to read off the incoming datagram and then fork and exit, or to process the arriving datagram and then time out.

In starting an rcp server from inetd, the only differences from the usual code are that svcudp_create should be called as:

```
transp = scvudp_create(0)
```

Since inetd passes a socket file as descriptor 0, and svc_register should be called as:

```
svc_register(PROGNUM, VERSNUM, service, transp, 0)
```

With the final flag as 0, since the program will already have been registered by inetd. If you want to exit from the server process and return control to inetd, you must explicitly exit since scv_run never returns.

The format of entries in /etc/servers for rpc services is:

```
rcp udp server-program program-number version-number
```

Where *server-program* is the C code implementing the server, and *program-number*, *version-number* are the program, and version numbers, respectively of the service. The keyword udp is replaced by tcp for TCP-based services.

If the same program handles multiple versions, the version number can be specified as a range, for example, 1-2.

Inetd provides several "trivial" services internally: echo, discard, chargen (character generator), daytime (human readable time), and time (machine readable time in the form of the number of seconds since midnight, January 1, 1900). These services are all tcp based.

## CONFIGURATION FILE

Each field of the configuration file must contain an entry; entries are separated by tabs or spaces. Comments are lines starting with a pound sign (#). The fields are as follows:

service name (for rpc services, specify service name and version)
: name of a valid service in the file /etc/services. For internal services, the service name must be the official name of the service (that is, the first entry in /etc/services).

  For rpc services, the entry is the name of a valid service in the file /etc/rpc, a slash, and a version number of the form 1-*n*, for example, mountd/1-2.

socket type
: one of the following: stream, dgram (datagram), raw, rdm (reliably delivered message), or seqpacket (sequenced packet socket).

protocol
: valid protocol as specified in /etc/protocols, for example, tcp or udp.

  For rpc services, the protocol name is preceded by rpc and a slash, for example, rpc/tcp or rpc/udp.

wait/nowait
: applies only to datagram, rpc sockets. Other sockets must specify nowait.

  If a datagram server connects to its peer, freeing the socket so inetd can receive further messages on the socket, the server is said to be multi-threaded, and nowait should be specified. Datagram servers that process all incoming datagrams on a socket and eventually time out are said to be single-threaded, and wait should be specified. Comsat (biff) and talk are examples of single-threaded datagram servers. Tftpd is an exception; it is a datagram server that establishes pseudo-connections. It must be listed as wait to avoid a race condition; the server reads the first packet, creates a new socket, then forks and exits to allow inetd to check for new service requests to spawn new servers.

uid
: the user ID of the user that should own the server (to permit servers to run without root privileges).

server program
: pathname of the program to be executed by inetd when a request is found on its socket. If inetd provides the service internally, specify this entry as internal.

server program arguments
: just as they are normally starting with argv[0], the name of the program. If inetd provides the service internally, specify argv[0] as internal.

Inetd rereads its configuration file when it receives a hangup signal (SIGHUP). Services can be added, deleted, and modified when the configuration file is reread.

## FILES

| | |
|---|---|
| /etc/inetd.conf | configuration file |
| /etc/services | |
| /etc/protocols | |

## SEE ALSO

comsat, ftpd, rexecd, mountd, rlogind, rshd, services, telnetd tftpd rwalld
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: inetd.conf

# init

## system initialization and process control

Performs system initialization for Diagnostix-2800 and Concentrix-2800. Starts Concentrix-2800 in single-user or multi-user mode depending on how the system was started or restarted. In multi-user mode, creates a process for each terminal port on which a login can occur. Terminates multi-user mode when appropriately signaled.

Note: Not intended for user invocation.

---

## /etc/init

---

### DESCRIPTION

Init is invoked by Diagnostix-2800 and then Concentrix-2800 at system start-up:

- **Diagnostix-2800** – Init performs system initialization for Diagnostix-2800 and executes the script rc.diag (which invokes console_loop).
- **Concentrix-2800** – Init performs system initialization for Concentrix-2800 and executes the script rc (which invokes rc.local). If reboot –s is in effect or rc fails, Concentrix-2800 stops in single-user mode by giving the superuser a shell on the console. By default, the single-user shell is the Bourne shell (/bin/sh). You can specify a shell by creating a symbolic link to a shell and naming the link /bin/rshell, for example, ln –s /bin/csh /bin/rshell makes the C shell the single-user shell. Type CTRL D to to exit from single-user mode and resume system startup.

In multi-user mode, init reads the file /etc/ttys and executes a command (usually /etc/getty) for each terminal specified in the file. Init calls getty with an argument as specified by the ttys file line. Getty reads the user's name and invokes login to log in a user and execute a shell.

When a login shell terminates (due to an end-of-file either typed explicitly or generated by a hang-up), init removes the appropriate entry from the file /etc/utmp, which records current users, and makes an entry in /usr/adm/wtmp, which maintains a history of logins and logouts; the wtmp entry is made only if a user logged in successfully on the line. Then the appropriate terminal is reopened and getty is reinvoked.

You can drop or add phone lines without rebooting the system by changing the file /etc/ttys and sending a hangup signal to the init process: type kill –HUP 1. Init interprets the hangup signal to mean that the file /etc/ttys should be read again. The shell process on each line that used to be active in /etc/ttys but is no longer there is terminated; a new process is created for each added line; lines unchanged in the file are undisturbed.

You can shut down the system to single-user mode by sending a terminate signal to the init process: type kill –TERM 1. Init does not wait for outstanding processes that are deadlocked (due to hardware or software failure) to die (which might take forever), but times out after 30 seconds and prints a warning message.

You can have the system slowly die away by sending a terminate stop signal to the init process: type kill -TSTP 1. Init ceases creating new getty processes. A later hangup signal resumes full multi-user operations; a later terminate signal shuts down the system to single-user mode.

If init dies, the system reboots itself automatically. If, at bootstrap time, init cannot be located, the system loops in user mode.

## DIAGNOSTICS

- `/etc/getty getty gettyargs failing, sleeping.`

  A process being started to service a line exits quickly each time it is started. Often caused by a ringing or noisy terminal line. Init sleeps for 30 seconds.

- `WARNING: Something is hung (wont die); ps axl advised.`

  A process is hung and could not be killed when the system was shutting down. This condition is usually caused by a process that is stuck in a device driver due to a persistent device error condition.

- `attempted execution by user`

  Only the Concentrix-2800 or Diagnostix-2800 kernel can invoke init.

## FILES

| | |
|---|---|
| /dev/console | console device |
| /dev/tty | terminal device |
| /dunix | Diagnostix kernel image |
| /etc/fstab | static information about filesystems |
| /etc/gettytab | terminal configuration data base |
| /etc/nologin | empty file whose existence prevents logins |
| /etc/rc | command script for starting Concentrix-2800 |
| /etc/rc.diag | command script for starting Diagnostix |
| /etc/ttys | terminal initialization data |
| /etc/utmp | login records (all logins and logouts) |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /usr/adm/wtmp | login record of all logins and logouts |
| /vmunix | a soft link to /unix |

## SEE ALSO

fastboot, getty, rc, reboot, shutdown
*Concentrix-2800 Commands and Applications Manual*: login, kill
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ttys

# INSTALL

## install diskless/NFS architecture and clients

Guides the superuser through installation procedures for a server of diskless/NFS clients. Allows the user to select and install client architectural support on the server from the client manufacturer's software distribution tape, and to create the environment on the server for specific client workstations. INSTALL asks the user to supply various parameters to be used during the installation process. Suitable defaults are provided for most prompts. Notes in this manual page should be read before running the script.

---

## /usr/etc/install/INSTALL

---

Before starting, you should take a few minutes to consider the server disk space requirements and layout. You need room for these trees:

- **roots** – Client root trees, one for each diskless/NFS client served. Initial space used is approximately 1.3 megabytes per client, but room must be left for tmp and spool space for each client.

- **swaps** – Client swap files, one for each diskless/NFS client served. A good initial value is 16 megabytes per client.

- **dumps** – Client crash dump files. These are initially created empty, so no space is used at first. Dump files are used when a client workstation crashes and a memory dump is saved by the savecore program. If no dump file is specified for a client, savecore instead uses the swap file for that client.

- **execs** – Architecture-dependent /usr filesystem used by the diskless/NFS clients. Space required is approximately 30–50 megabytes per architecture type, depending on optional software selected.

- **home** – A general work area set up to be mounted by all diskless/NFS clients. The subdirectory /home/*servername* is created for each server. The intent is that the indivual user home directories are located in these trees, but other common (local) software may be placed there as well. The home path may be set to none if this setup is not desired.

If individual diskless/NFS clients are to be installed when installing architectural support, the administrator must have the following information about each client before running INSTALL: architecture type, ethernet address, and IP address. The network addresses must already be entered into the appropriate administration files before running INSTALL.

You can use INSTALL to add new diskless/NFS clients after the base architecture is installed. The script detects this fact and presents the administrator with several options, as described below.

## PARAMETER VALUES

The script collects the following information:

- **tape drive** – The tape drive can either be local or remote. If a remote tape drive is selected, you are prompted for the machine name of the remote host. The machine name is verified in the network information service (via ypmatch), or in /etc/hosts if NIS is not running on the server (detected by checking the domain name returned by the domainname command). You should set the domain name to noname or leave it NULL if you do not want the NIS to be used during diskless/NFS client installation steps.

- **tape type** – Once the tape location is specified, you are asked for the tape type. The tape type can be one of the default driver designations listed, or you can specify the full pathname of a non-rewinding raw tape device (such as /dev/nrmt0). Note that only one tape device can be specified per invocation of INSTALL. If you have different tape formats for different architectures, you must run INSTALL separately for each tape type.

- **client architectures** – The next step is to specify the various client architectures being loaded during this process. The architecture names used should be the same as the designations on the software distribution tapes from the client workstation manufacturer. For each architecture entered, a pathname is requested where the software is to be loaded on the server. The architecture and release name are automatically appended to the entered path. A sample response to this prompt is /export/exec; for a sun3 distribution tape for SunOS 4.0, this response makes the actual install location for the software /export/exec/sun3. Once the pathname is entered for one architecture, it can be used as a default for following architectures by simply pressing RETURN to the prompts.

- **executables** – After the exec pathname is specified, the script determines if the executables for that architecture are already loaded. If so, you can elect to ignore the entered architecture, remote the existing tree and reload it, load additional packages for the architecture, or leave the executable tree alone and proceed. In all cases but the first, you are allowed to enter information for adding new clients of that architecture.

- **software selections** – When selecting new software to load, the tape's table of contents is searched to obtain the software choices available for loading from that distribution. You must select options to load by answering yes or no to the prompts as they are presented. The prompts include the size of the particular software option, together with an indication as to the importance of loading the option. A load list is constructed in /tmp/EXTRACTLIST.*architecture* to be used in setup_exec.

- **shared objects** – Once the software selections are made for a given architecture, you are asked whether shared objects should be shared by clients of different architectures. Shared objects refer to such items as manual pages and common library support packages, such as timezone information and system V terminfo databases. The benefit of sharing is that disk space is saved when two or more architectures are loaded onto the same server. Sharing is not necessary (and should not be selected) if the server is to support only one diskless/NFS client architecture.

- **clients** – When all desired architectures are specified, enter done to the next architecture prompt. At this point, you are asked to specifiy the individual clients for each architecture. It is required that the ethernet addresses (/etc/ethers or the ethers NIS map) and IP addresses (/etc/hosts or the hosts NIS map) are updated with correct information for all new clients being added. This information is verified during the installation process similarly to the remote tape hostname mentioned above. Note that the special character ^ can be used as input to any prompt requesting client information to cause that client information to be discarded, and a new client can be specified. This feature can be used to correct mistakes after they are entered (but before reaching the end of specification requests for that client). Also, information entered in a particular category for one client can be used for following clients by entering only a RETURN in response to the following prompts.

  Additional information required for clients is the NIS type (master, slave, client, or none). Most likely choices here are client or none. Swap size for the client can be specified in a number of different units. Enter anything but a digit at the prompt for examples. A typical value is 16M (for 16 megabytes).

  The various pathnames that are specified (for root, swap, dump) are the locations on the server where these functions are installed. The entered pathname should be the generic path; the client's hostname is appended to make the actual location. A sample input pathname is /export/root (which turns into /export/root/*clientname* during that client's installation. The home path is where the user home directories are installed. The convention is that users on workstations being served by the server have home directories on /home/*servername/username*. A typical answer to this prompt is /home. Defaults provided in the script for root, swap, dump, and home are /export/root, /export/swap, none, and /home, respectively.

When all client information for a given architecture is complete, enter done to the prompt for the next client name. When all architectures are complete, you are asked to proceed. If you enter no to this question, the current state is left in the configuration files in /tmp. You can continue from this point by rerunning INSTALL and entering continue to the first prompt asking for architecture names. When installation begins, each architecture and its clients are installed in turn. The script setup_exec is first invoked to load the software from the distribution tapes. Then setup_client is invoked for each client of that architecture.

## DIAGNOSTICS

Self-explanatory.

## FILES

/tmp/*.<architecture>

## SEE ALSO

opt_software, setup_client, setup_exec
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ethers, hosts

# keyenvoy

## talk to keyserver

Keyenvoy is used by some RPC programs to talk to the key server (keyserv). The key server only talks to a root process; keyenvoy is a set-uid root process that acts as an intermediary between a user process that wishes to talk to the key server and the key server itself.

This program cannot be run interactively.

---

**keyenvoy**

---

### SEE ALSO

keyserv

# keyserv

server for storing public and private keys

Keyserv is a daemon that stores the private encryption keys of each user logged into the system. These encryption keys give access to secure network services such as secure NFS.

---

**keyserv [-n ]**

---

## ARGUMENTS

**-n** Do not read root's key from /etc/.rootkey. Instead, prompt the user for the password to decrypt root's key stored in the NIS and then store the decrypted key in /etc/.rootkey for future use. This option is useful if the /etc/.rootkey file ever gets out of date or corrupted.

## DESCRIPTION

When a user logs in to the system, the login program uses the login password to decrypt the user's encryption key stored in the NIS and then gives the decrypted key to the keyserv daemon to store away.

Normally, root's key is read from the file /etc/.rootkey when the daemon starts up. This is useful during power-fail reboots when no one is around to type a password, yet you still want the secure network services to operate normally.

## FILES

/etc/.rootkey

## SEE ALSO

*Concentrix Commands and Applications:* login
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* publickey

# kgmon

## control kernel profiling

Controls the collection and dumping of kernel profiling data, including processor states and program counter histograms.

---

**/etc/kgmon [-b[*cpuid#*] -h -r -p -c -u *system-memory*]**

---

### ARGUMENTS

**-b**[*cpuid#*]
: begins collecting profile data for processor *cpuid#*; defaults to all processors.

**-h**
: halts collecting profile data.

**-r**
: resets data collection buffers in the kernel.

**-p**
: produces a profile data file, kgmon.out, for input to the kernel profile report program, kgprof.

**-c**
: prompts for user comments to be added to the profile data file.

**-u**
: deletes the lock file.

*system-memory*
: specifies the name of the system image file that contains the kernel symbols; defaults to /vmunix.

### DESCRIPTION

The flag enbproflag must be set to a non-zero value during system bootstrap to enable kernel profiling (more specifically, to direct the kernel to allocate the necessary profile buffers). Kernel profiling is single threaded through the lock file /tmp/kgproflock; the owning uid is the holder of the lock.

### EXAMPLE

The following sequence of commands illustrates the use of kgmon:

```
% kgmon -r              | Reset the kernel profiling buffers
% kgmon -b              | Start profiling all cpus
commands used to test the system
% kgmon -h -p -c        | Stop profiling and dump the profile
                        | buffers to kgmon.out with a comment
```

## kgmon.out FORMAT

kgmon header (/usr/include/kgmon.h)
low pc, high pc, and number of pc historgram cells
hz value
mask of profiled processors
processor state counters (/usr/include/sys/kgprof.h)
array of pc histogram cells

## FILES

| | |
|---|---|
| kgmon.out | profile data file |
| /vmunix | system image file |
| /tmp/kgproflock | profiling lock file |
| /dev/kmem | system memory access file |

## SEE ALSO

kgprof

# kgprof

## display kernel execution profile

Reports kernel profiling data, including processor states and program counter histograms. The report is produced on standard output.

---

**/usr/ucb/kgprof** [*option* ...] [*kernel-image* [*kgmon.out* ...]]

---

### ARGUMENTS

**-a** suppresses the display of statically declared functions. If this option is specified, all relevant information about the static function (for example, time samples, calls to other functions, calls from other functions) belongs to the function loaded just before the static function in the kernel image file.

**-b** supresses the display of a description of each field in the profile.

**-c** the static call graph of the program is discovered by a heuristic that examines the text space of the object file. Static-only parents or children are indicated with call counts of 0.

**-e** *name*
suppresses the display of the graph profile entry for *name* and all its descendants (unless they have other ancestors that are not suppressed). More than one **-e** option can be specified. Only one name can be specified with each **-e** option.

**-E** *name*
suppresses the display of the graph profile entry for *name* (and its descendants) as in **-e** above, and also excludes the time spent in name (and its descendants) from the total and percentage time computations. (For example, **-E mcount -E mcleanup** is the default.)

**-f** *name*
displays the graph profile entry of only *name* and its descendants. More than one **-f** option can be specified. Only one name can be specified with each **-f** option.

**-F** name
displays the graph profile entry of only *name* (and its descendants) as in **-f**, above, and also uses only the times of the displayed routines in total time and percentage computations. More than one **-F** option can be specified. Only one name can be specified with each **-F** option. The **-F** option overrides the **-E** option.

**-s** produces or adds to a file kgmon.sum, which represents the sum of the profile information in all the specified profile files.

**-w** sorts the lock statistics by lock hold time. The default is by number of acquisitions.

**-z** displays routines with zero usage.

## DESCRIPTION

The effect of called routines is incorporated in the profile of each caller. The profile data is taken from the kernel profile file (kgmon.out by default) created by the kgmon utility.

The symbol table in the named object file (/vmunix by default) is correlated with the call graph profile file. If more than one profile file is specified, the kgprof output shows the sum of the profile information in the given profile files.

First, a profile of the system processors is given including time spent in system, user, idle, and lock spin state.

Next, a lock profile is given for all system multiprocessor locks including the number of acquisitions, collisions, and average wait times for each lock. In addition, if a lock trace was performed by kgmon, a detailed trace of lock hold times is provided by call site.

Next, a flat profile is given, similar to that provided by gprof. This listing gives the total execution times and call counts for each of the functions in the kernel, sorted by decreasing time.

Next, these times are propagated along the edges of the call graph. Cycles are discovered, and calls into a cycle are made to share the time of the cycle. A second listing shows the functions sorted according to the time they represent, including the time of their call graph descendents. Below each function entry is shown its (direct) call graph children, and how their times are propagated to this function. A similar display above the function shows how this function's time and the time of its descendents is propagated to its (direct) call graph parents.

Cycles are also shown, with an entry for the cycle as a whole and a listing of the members of the cycle and their contributions to the time and call counts of the cycle.

## FILES

| | |
|---|---|
| /vmunix | kernel image |
| kgmon.out | kernel profile raw data file |
| kgmon.sum | summarized profile data file |

## SEE ALSO

kgmon
*Concentrix Commands and Applications*: gprof

## NOTES

Beware of quantization errors. The granularity of the sampling is shown, but remains statistical at best. We assume that the time for each execution of a function can be expressed by the total time for the function divided by the number of times the function is called. Thus the time propagated along the call graph arcs to parents of that function is directly proportional to the number of times that arc is traversed.

Parents that are not themselves profiled have the time of their profiled children propagated to them, but they appear to be spontaneously invoked in the call graph listing, and do not have their time propagated further. Similarly, signal catchers, even though profiled, appear to be spontaneous (although for more obscure reasons). Any profiled children of signal catchers should have their times propagated properly, unless the signal catcher was invoked during the execution of the profiling routine, in which case all is lost.

# lockd

## network lock daemon

Processes lock requests that are either sent locally by the kernel or remotely by another lock daemon.

The lockd daemon is typically invoked by /etc/rc.local. (As distributed, the lockd invocation is commented out in /etc/rc.local. Uncomment the invocation if you want lockd to start automatically with the system.)

---

**/etc/rpc.lockd [-t *timeout*] [-g *graceperiod*]**

---

### ARGUMENTS

-t *timeout* interval in seconds to retransmit a lock request to a remote server; defaults to 15 seconds.

-g *graceperiod*
the grace period duration in seconds; defaults to 45 seconds.

### DESCRIPTION

Lockd forwards lock requests for remote data to the server site's lock daemon through the RPC/XDR package. Lockd then requests the status monitor daemon, statd, for monitor service. The reply to the lock request is not sent to the kernel until the status daemon and the server site's lock daemon have replied.

If either the status monitor or server site's lock daemon is unavailable, the reply to a lock request for remote data is delayed until all daemons become available.

When a server recovers, it waits for a grace period for all client site lock daemons to submit reclaim requests. Client site lock daemons, on the other hand, are notified by the status daemon of the server recovery and promptly resubmit previously granted lock requests. If a lock daemon fails to secure a previously granted lock at the server site, the lock daemon sends SIGLOST to a process.

### SEE ALSO

statd
*Concentrix-2800 System Reference Manual: System Calls and C Library Routines*: fcntl, lockf, signal

# lpc

## line printer control program

Controls the operation of the line printer system. For each line printer configured in /etc/printcap, lpc can be used to:

- disable or enable a printer.
- disable or enable a printer's spooling queue.
- rearrange the order of jobs in a spooling queue.
- find the status of printers, and their associated spooling queues and printer dameons.

---

**/etc/lpc** [*command* [*argument* ...]]

---

### ARGUMENTS

With no arguments, lpc prompts for commands from standard input. If arguments are supplied, lpc interprets the first argument as a command (which can be abreviated) and the remaining arguments as parameters to the command. Standard input can be redirected to read commands from a file.

The following is the list of recognized commands. A vertical bar (|) indicates a choice of arguments.

**?** [ *command* ... ]
prints a short description of each command specified in the argument list, or, if no arguments are given, a list of the recognized commands. Use a backslash to escape the question mark.

**abort** { **all** | *printer* ... }
terminates an active spooling daemon on the local host immediately and then disables printing (preventing new daemons from being started by lpr) for the specified printers.

**clean** { **all** | *printer* ... }
removes any temporary files, data files, and control files that cannot be printed (that is, that do not form a complete printer job) from the specified print queues on the local machine.

**disable** { **all** | *printer* ... }
prevents new jobs from being entered into the specified print queues.

**down** { **all** | *printer* } *message*
turns the specified print queues off, disables printing, and puts the string *message* in the printer status file. *Message* need not be quoted. The command lpq, if invoked, prints *message*.

**enable** { **all** | *printer* ... }
enables new jobs to be entered into the specified queues.

**exit** exits from lpc.

**help** [ *command* ... ]
prints a short description of each command specified in the argument list, or, if no arguments are given, a list of the recognized commands.

**quit** exits from lpc.

**restart** { **all** | *printer* ... }
attempts to start a new printer daemon. Useful when some abnormal condition causes the daemon to die unexpectedly leaving jobs in the queue (lpq reports that no daemon is present).

**start** { **all** | *printer* ... }
enables printing and starts a spooling daemon for the listed printers.

**status** [ *printer* ... ]
displays the status of daemons and queues on the local machine.

**stop** { **all** | *printer* ... }
stops a spooling daemon after the current job completes and disables printing.

**topq** *printer* [ *jobnum* ... ] [ *user* ... ]
places the jobs at the top of the printer queue in the order listed.

**up** { **all** | *printer* }
enables queueing and printing for the specified print queue, and starts a spooling daemon.

## DIAGNOSTICS

| | |
|---|---|
| ?Ambiguous command | abreviation matches more than one command |
| ?Invalid command | no match was found |
| ?Privileged command | command can be executed by root only |

## FILES

| | |
|---|---|
| /etc/printcap | printer description file |
| /usr/spool/* | spool directories |
| /usr/spool/*/lock | lock file for queue control |

## SEE ALSO

lpd
*Concentrix Commands and Applications Manual*: lpr, lpq, lprm
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: printcap

# lpd

## line printer daemon

Makes a single pass through /etc/printcap to set up the existing printers and print any files left after a crash. Uses the system calls listen and accept to receive requests to print files in the queue, transfer files to the spooling area, display the queue, and remove jobs from the queue. For each request, lpd forks a child to handle the request so that the parent can continue to listen for more requests.

The lpd daemon is typically invoked at system startup by /etc/rc.

---

## /usr/lib/lpd [-l] [*port*]

---

### ARGUMENTS

-l logs valid requests received from the network. Useful for debugging purposes.

*port* specifies the Internet port number used to rendezvous with other processes (normally obtained with getservbyname).

### DESCRIPTION

Remote access control is provided by two means. First, all requests must come from one of the machines listed in the file /etc/hosts.equiv or /etc/hosts.lpd. Second, if the rs capability is specified in the /etc/printcap entry for the printer being accessed, lpr requests are only honored for those users with accounts on the machine with the printer.

The file minfree in each spool directory contains the number of disk blocks to leave free so that the line printer queue does not completely fill the disk. You can edit this file.

Lpd uses flock to provide exclusive access to the file lock in each spool directory and to prevent multiple deamons from becoming active simultaneously. If the daemon should be killed or die unexpectedly, the lock file need not be removed. The lock file is kept in a readable ASCII form and contains two lines. The first is the process id of the daemon and the second is the control file name of the current job being printed. The second line is updated to reflect the current status of lpd for the programs lpq and lprm.

After lpd successfully sets the lock, it scans the spooling directory for files whose names begin with cf. Each line in a cf file specifies information or an action to be taken. The first character of each line is a key that identifies the remainder of the line:

C Classification. String to be used for the classification line on the burst page.

c Cifplot file. The file contains data produced by cifplot.

d DVI file. The file contains tex output.

**f** Formatted file. Name of a file to print that is already formatted.

**g** Graph file. The file contains data produced by plot.

**H** Host Name. Name of the machine where lpr was invoked.

**I** Indent. The number of characters to indent the output (in ASCII).

**J** Job Name. String to be used for the job name on the burst page.

**L** Literal. Identification information from the password file. Causes the banner page to be printed.

**l** Like f but passes control characters and does not make page breaks.

**M** Mail. Sends mail to the specified user when the current print job completes.

**N** File name. The name of the file that is being printed, or a blank for the standard input (when lpr is invoked in a pipeline).

**n** Ditroff file. The file contains device independent troff output.

**P** Person. Login name of the person who invoked lpr. Used to verify ownership by lprm.

**p** Pr. Name of a file to print using pr as a filter.

**r** The file contains text data with Fortran carriage control characters.

**T** Title. String to be used as the title for pr.

**t** Troff file. The file contains troff output (cat phototypesetter commands).

**U** Unlink. Name of file to remove upon completion of printing.

**v** The file contains a raster image.

**W** Width. Changes the page width (in characters) used by pr and the text filters.

**1** Troff Font R. Name of the font file to use instead of the default.

**2** Troff Font I. Name of the font file to use instead of the default.

**3** Troff Font B. Name of the font file to use instead of the default.

**4** Troff Font S. Name of the font file to use instead of the default.

If a file can not be opened, a message is placed in the log file (normally the console). Lpd tries up to 20 times to open a file it expects to be there, after which it skips the file to be printed.

In normal operation, lpd listens for service requests on multiple sockets, one in the UNIX domain (named /dev/printer) for local requests, and one in the Internet domain (under the printer service specification) for requests for printer access from other systems; see the system call socket and the file format *services* for more information.

Clients communicate with lpd using a simple transaction oriented protocol. Authentication of remote clients is based on the "privilege port" scheme employed by /etc/rshd (remote shell server) and the library subroutine rcmd. The following table shows the requests understood by lpd. In each request the first byte indicates the meaning of the request and is followed by the name of the printer to which it should be applied. Additional qualifiers can follow, depending on the request.

^A*printer*\n
: check the queue for jobs and print any found

^B*printer*\n
: receive and queue a job from another machine

^C*printer* [*users* ...] [*jobs* ...]\n
: return short list of current queue state

^D*printer* [*users* ...] [*jobs* ...]\n
: return long list of current queue state

^E*printer person* [*users* ...] [*jobs* ...]\n
: remove jobs from a queue

## DIAGNOSTICS

The lpd program can write many different messages to the error log file (the file specified in the lf entry in /etc/printcap). Most of these messages are about files that cannot be opened and usually indicate that the protection modes of the files are not correct. Files can also be inaccessible if people manually manipulate the line printer system (that is, bypass the lpr program).

## FILES

| | |
|---|---|
| /dev/printer | socket for local print requests |
| /dev/* | print devices |
| /etc/printcap | printer capability data base |
| /etc/hosts.equiv | list of trusted hosts |
| /usr/spool/lpd/* | spooling area (queue) for print jobs |

## SEE ALSO

lpc
*Concentrix Commands and Applications Manual*: pac, lpr, lpq, lprm, pr, troff
*Concentrix-2800 System Reference Manual: System Calls and C Library Routines*: listen, accept, flock, getservbyname, plot, syslog
*Concentrix-2800 System Reference Manual: Miscellaneous Routines, Special Files, and File Formats*: printcap

# mailq

## print contents of queue

Prints the contents of the mail queue. Identical to sendmail -bp.

---

## /usr/ucb/mailq

---

### DESCRIPTION

Each entry of the display contains the following information:

- **Queue identification** – Internal identification value for the queue entry.
- **Size** – Size of the message in bytes.
- **Time** – Time that the message arrived in the mail queue.
- **Sender** – The user name of the sender of the message.
- **Status** – Any status status information on the message.
- **Recipient** – The user name of the intended recipient.

### EXAMPLE

The following example demonstrates a mailq display:

```
% /usr/ucb/mailq
                Mail Queue (13 requests)
--QID-- --Size-- -----Q-Time----- ------------Sender/Recipient------------
AA01453       236 Sun Jan  3 14:48 jake
                                   (Deferred: Connection timed out)
                                   tom@early
AA02365      1393 Sun Jan  3 15:31 <julian@abel>
                                   (Deferred: Connection timed out)
                                   jake@early
AA00826       434 Sun Jan  3 13:28 bob
                                   (Deferred: Connection timed out)
                                   larry@early
                                   lance@early
AA08674       111 Sun Jan  3 00:00 root
                                   (Deferred: Connection timed out)
                                   george@early
AA08642       111 Sun Jan  3 00:00 root
                                   (Deferred: Connection timed out)
                                   george@early
```

### FILES

/usr/spool/mqueue/* sendmail queue and system log

### SEE ALSO

sendmail

# mailstats

## print sendmail statistics

Mailstats prints out the statistics collected by the sendmail program on mailer usage.

---

**/usr/etc/mailstats** [*filename*]

---

### DESCRIPTION

Sendmail collects these statistics if the file indicated by the S*file* configuration option of sendmail exists.

The mailstats program first prints the time that the statistics file was created and the last time it was modified. It then prints a table with one row for each mailer specified in the configuration file. The first column is the mailer number, followed by the symbolic name of the mailer. The next two columns refer to the number of messages received by sendmail, and the last two columns refer to messages sent by sendmail. The number of messages and their total size (in 1024 byte units) is given. No numbers are printed if no messages were sent (or received) for any mailer.

You might want to add an entry to /var/spool/cron/crontab/root to reinitialize the statistics file once a night. Copy /dev/null into the statistics file or otherwise truncate it to reset the counters.

### NOTES

- Mailstats should read the configuration file instead of having a hard-wired table mapping mailer numbers to names.

### FILES

| | |
|---|---|
| /etc/sendmail.st | default statistics file |
| /etc/sendmail.cf | sendmail configuration file |
| /var/spool/cron/crontab/root | |

### SEE ALSO

sendmail

# makedbm

make a network information service dbm file

Builds dbm files for the NIS data base.

---

**/etc/yp/makedbm [-b][-i *name*][-o *name*][-d *name*][-m *name*] *infile outfile***

**/etc/yp/makedbm [-u *name*]**

---

## ARGUMENTS

**-b** Interdomain. Propagate a map to all servers using the interdomain name server named.

**-i *name*** creates a special entry with the key yp_input_file.

**-o *name*** creates a special entry with the key yp_output_name.

**-d *name*** creates a special entry with the key yp_domain_name.

**-m *name*** creates a special entry with the key yp_master_name. Defaults to the local host name.

**-u *name*** (undo) prints out a dbm file one entry per line, with a single space separating keys from values.

***infile*** specifies an input file. A hyphen specifies the standard input.

***outfile*** specifies the base name of the output files.

## DESCRIPTION

Converts *infile* to a pair of files in dbm format, namely *outfile*.pag and *outfile*.dir. Each line of the input file is converted to a single dbm record. All characters up to the first tab or space form the key and the rest of the line is the data. If a line ends with a backslash (\), the data for that record is continued on to the next line. Clients can interpret the pound sign (#) as desired; makedbm does not treat it as a comment character. *infile* can be '-', in which case the standard input is read.

Makedbm generates a special entry with the key yp_last_modified, which is the date of *infile* (or the current time if *infile* is -).

## EXAMPLE

The following shell script converts a standard /etc/passwd file to the key-value form used by makedbm. That is, the key is a username, and the value is the remaining line in the /etc/passwd file.

```
#!/bin/awk -f
BEGIN { FS = ":"; OFS = "\t"; }
{ print $1, $0 }
```

Look at /etc/yp/Makefile for additional examples.

## SEE ALSO

named. ypinit, ypmake

*Concentrix Commands and Applications:* yppasswd
*Concentrix-2800 System Reference Manual*: *System Calls and C Library Routines*: ndbm

# make_dirs

make a directory and its parents

---

**/usr/etc/install/make_dirs** *dirname* [*mode*]

---

Makes the directory *dirname* and all necessary parent directories leading to it. If *mode* is specified, all created directories are set to that mode; otherwise, the newly created directories have modes subject to the user's umask value.

**SEE ALSO**

chmod, mkdir, umask

# makefs

## construct a file system

Constructs a file system with a root directory and a lost+found directory on a specified disk partition or disk stripe. Does not initialize a boot program. The number of i-nodes is calculated as a function of the file system size.

---

**/etc/makefs** [*option ...*] *special*

---

### ARGUMENTS

Each value argument (except *type*) is a decimal integer. If the decimal integer is suffixed by k or K, the value is multiplied by 1024.

**-a** *n* specifies the maximum number of contiguous blocks laid out before forcing a rotational delay (see -d below). The default value is 1024, since most device drivers require an interrupt per disk transfer. Device drivers that can chain several buffers together in a single transfer should set this to the maximum chain length.

**-b** *blksize* specifies the primary blocksize of the file system in bytes. A 4.3BSD file system must be 4096. A track-allocated file system must be a multiple of 4096 and defaults to one track of the disk (or less to make a multiple of 4096).

**-B** specifies size of partition in sectors: overrides ioctl of special file.

**-c** *ncpg* specifies the number of disk cylinders per cylinder group. The default is 24 for track-allocated file systems and 16 for 4.3BSD file systems.

**-d** *n* specifies the expected time (in milliseconds) to service a transfer completion interrupt and initiate a new transfer on the same disk. It is used to decide how much rotational spacing to place between successive blocks in a file. The default delay is set to 5. Because of the -a default value of 1024, this only matters when blocks cannot be allocated sequentially.

**-e** *n* indicates the maximum number of blocks any single file can allocate out of a cylinder group before it is forced to begin allocating blocks from another cylinder group. The default value is 1024 unless 1024 times the block size is more than a quarter of a cylinder group. In this case, the default value is halved to 512. If 512 times the block size is more than a quarter of a cylinder group, the default value is again halved to 256.

**-f** *fragsize* specifies the fragment size of the file system (the smallest amount of disk space that can be allocated to a file) in bytes. A 4.3BSD file system must be 512, 1024 (default), 2048, or 4096. A track-allocated file system must be a power of two that is larger than 2048 (defaults to 4096).

**-F** *fullblk* (track-allocated file systems only) specifies the maximum size of a file for which fragments can be allocated. For example, if you specify 7, the last

block of a seven-block file can be a fragment; the last block of an eight-block file cannot. The value must be in the range 0 to 12. The default is 1.

**-i** *indirsize* (track-allocated file systems only) specifies the size of an indirect pointer block. The size must be a power of two not smaller than *fragsize* and not larger than *blksize*. The default is *fragsize*.

**-I** *nbpi* specifies the number of bytes of file system per i-node. The default is the greater of 2048 and *fragsize*.

**-m** *minfree* specifies the minimum percentage of free disk space allowed. Once the file system capacity reaches this threshold, only the super-user is allowed to allocate disk blocks. The default is ten percent.

**-n** describes the file system that would be created instead of creating it.

**-S** specifies number of sectors per track: overrides ioctl.

**-t** *type* specifies the type of file system:

**4.3** 4.3BSD file system (4.2 acts the same as 4.3)

**tfs** track-allocated file system

If the type of file system is not specified, the fragment size determines the type of file system. Values smaller than 4096 create a 4.3BSD file system; larger values create a track-allocated file system. Omitting both the -f and -t options creates a track-allocated file system.

**-T** specifies number of tracks per cylinder: overrides ioctl.

*special* specifies the raw device name of the disk partition or disk stripe on which to write the file system.

## NOTES

- One-gigabyte drives are formatted at 80 sectors per track on the FX/2800 system.
- The file system does not use partial blocks. Any leftover fragments at the end of the file system, before the end of the partition, are discarded.

## DIAGNOSTICS

- `Not a typewriter` – Usually caused by specifying the block device instead of the raw device.

## EXAMPLES

The following example makes a 4.3BSD file system on /dev/rxdv01h:

```
# /etc/makefs -t 4.3 /dev/rxdv01h
```

The following example makes a track-allocated file system on /dev/rxdv01h:

```
# /etc/makefs -t tfs /dev/rxdv01h
```

The following example makes a track-allocated file system on /dev/rstripe0:

```
# /etc/makefs -t tfs /dev/rstripe0
```

## SEE ALSO

dumpfs, fsck, mkstripe, tunefs
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: fs, dir

# MAKEHOSTS

## set up /usr/hosts directory

The MAKEHOSTS command sets up the directory /usr/hosts so that the rlogin and rsh commands can be invoked in the form of a remote host name. It creates a symbolic link to /usr/ucb/rsh for each remote host name in /etc/hosts or, if your system is running the NIS, for each host name returned by ypcat hosts.

---

## /usr/hosts/MAKEHOSTS

---

### DESCRIPTION

Once the directory /usr/hosts is set up, users who have it in their search paths can invoke the rlogin and rsh commands by specifying the name of a remote host. For example, if the directory /usr/hosts contains a file kryptn, then:

```
% kryptn
```

is equivalent to:

```
% rlogin kryptn
```

and:

```
% kryptn pwd
```

is equivalent to:

```
% rsh kryptn pwd
```

### FILES

/usr/hosts/[a-z]* remote host names

### SEE ALSO

*Concentrix Commands and Applications Manual*: rlogin, rsh

# makekey

## generate encryption key

Makekey improves the usefulness of encryption schemes depending on a key by increasing the amount of time required to search the key space. It reads 10 bytes from its standard input, and writes 13 bytes on its standard output. The output depends on the input in a way intended to be difficult to compute (that is, to require a substantial fraction of a second).

---

## /usr/lib/makekey

---

### DESCRIPTION

The first eight input bytes (the *input key*) can be arbitrary ASCII characters. The last two (the *salt*) are best chosen from the set of digits, upper- and lower-case letters, period and slash. The salt characters are repeated as the first two characters of the output. The remaining 11 output characters are chosen from the same set as the salt and constitute the *output key*.

The transformation performed is essentially the following: the salt is used to select one of 4096 cryptographic machines all based on the National Bureau of Standards DES algorithm, but modified in 4096 different ways. Using the input key as key, a constant string is fed into the machine and recirculated a number of times. The 64 bits that come out are distributed into the 66 useful key bits in the result.

Makekey is intended for programs that perform encryption (for instance, ed and crypt). Usually makekey's input and output are pipes.

### SEE ALSO

*Concentrix Commands and Applications Manual*: crypt, ed

# makeramdisk

## create RAM disk image

Creates an image file under Diagnostix-2800 that is used to build a RAM disk under multi-user Concentrix-2800. A RAM disk simulates a disk device in physical memory. The physical memory used for the RAM disk is dedicated, that is, it is not available for virtual memory.

**NOTE:** This utility runs under Diagnostix-2800.

---

## /console/makeramdisk *file size*

---

### ARGUMENTS

*file* is the name of the RAM disk image in standard file specification format. Also created is the file ramdisktab, in the current working directory, which is formatted like disktab and contains a single entry for the RAM disk.

*size* is the size of the RAM disk in whole megabytes (must be an integer). The RAM disk image is created with 1MB "cylinders" for numerical convenience; each "cylinder" has 4 "heads" and 256KB "tracks." The entire RAM disk is allocated to one r partition.

### DESCRIPTION

The full procedure for creating a RAM disk for use under multi-user Concentrix-2800 is as follows:

1. Shut down the system to (or start the system stopping in) Diagnostix-2800 mode. For example (from Concentrix-2800):

```
Concentrix# /etc/shutdown -h now Quick reboot to get RAM disk.
```

2. When you have the Diagnostix-2800 prompt, change directory to /console and execute makeramdisk. For example:

```
dunix# cd /console
dunix# makeramdisk ramdisk 2
0+1 records in
1+0 records out
0+1 records in
1+0 records out
Created 2 MB ramdisk root partition.
```

3. Create a file srbootflags that contains the line -r *file*, where *file* is the name of the RAM disk image. For example:

```
dunix# cat > srbootflags
-r ramdisk
CTRL d
```

4. Start Concentrix-2800, for example:

```
dunix# CTRL d
FX/2800 Startup

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help): y
```

During FX/2800 booter mode, messages of the following form appear:

```
Alliant FX/2800 Booter
Allocating 512 pages for ramdisk image of 512 pages.
Ramdisk of 512 pages loaded with 1024 bytes.
```

During single-user mode, a message of the following form appears:

```
ip00> ramdisk: "ramdisk" <1024.,4.,512.> spares=0
```

5. Log into Concentrix-2800 as root. If raw and block devices for the RAM disk (major number 0x89, minor number 0x11) do not exist in /dev, make them. If a mount point does not exist for the RAM disk device, use mkdir to create one in the root directory. (These are one-time procedures.) For example:

```
Concentrix# mknod /dev/rramdisk c 0x89 0x11
Concentrix# mknod /dev/ramdisk b 0x89 0x11
Concentrix# mkdir /ram
```

6. Make a file system (track allocated or 4.3BSD) for the RAM disk device; specify -b 128k for a track allocated file system. For example:

```
Concentrix# makefs -t 4.2 /dev/rramdisk
This ... 4.2 BSD filesystem of 4096 [512 byte] blocks on '/dev/rramdisk'
Each cylinder is 4 tracks of 262144 bytes.
The block size is 4096 and the fragment size is 1024.
There are 16 cylinders per group, with 2048 bytes per inode.
10% of the data blocks are reserved.
/dev/rramdisk:  4096 sectors in 2 cylinders of 4 tracks, 512 sectors
        2.1Mb in 1 cyl groups (16 c/g, 16.78Mb/g, 896 i/g)
super-block backups (for fsck -b#) at: 24,
```

7. Mount the file system at the mount point. For example:

```
Concentrix# mount /dev/ramdisk /ram
```

The RAM disk is ready for use; data and executable files can be copied to it from real disks.

> **WARNING:** If the system shuts down or crashes, everything on the RAM disk is lost.

Procedures 6 and 7 can be performed automatically at each system start-up. Place a line in /etc/rc *before* the mount commands to make the RAM disk file system. For example:

```
/etc/makefs -t 4.2 /dev/rramdisk
```

Place a line in /etc/fstab for the RAM disk file system, but specify 0 for the *passno* field (the last field). For example:

```
/dev/ramdisk   /ram    4.2     rw      1       0
```

## NOTES

- To use the RAM disk for /tmp: delete /tmp, make a tmp subdirectory in the RAM disk file system, and make /tmp a link to this directory. For example:

```
rm -r /tmp
mkdir /ram/tmp
ln -s ram/tmp /tmp
```

## DIAGNOSTICS

- `Badly formed number.` – The *size* specification must be an integer.

# makewhatis

## construct whatis data base

Constructs a new whatis data base in the current directory based on the contents a man data base. The whatis data base names the contents of the man data base and provides one-line descriptions of each entry.

---

## /usr/lib/makewhatis [*mandir*]

---

### ARGUMENTS

*mandir* top directory in a man data base; defaults to /usr/man.

### DESCRIPTION

The whatis data base should be built in or moved to the top directory of the man directory hierarchy, for example, /usr/man/whatis.

The script /usr/lib/makewhatis reconstructs the whatis data base. The script takes a single optional argument, the pathname of the man data base, which defaults to /usr/man. The man subdirectories that are processed are fixed in the script at 0 through 8, a, c, f, and s. The script employs the services of the program /usr/lib/getNAME to extract the whatis information from the man files.

### FILES

| | |
|---|---|
| /usr/lib/getNAME | program that extracts whatis line from man files |
| /usr/man/whatis | whatis data base |
| /usr/man/man?/* | on-line documentation in nroff source format |

### SEE ALSO

catman
*Concentrix Commands and Applications Manual*: apropos, man, nroff, whatis, whatisin

# mconnect

## connect to SMTP mail server socket

Mconnect opens a connection to the mail server on a given host, so that it can be tested independently of all other mail software. If no host is given, the connection is made to the local host. Servers expect to speak the Simple Mail Transfer Protocol (SMTP) on this connection. Exit by typing the quit command. Typing EOFsends an end of file to the server. An interrupt closes the connection immediately and exits.

---

**/usr/etc/mconnect** [ **-p** *port*] [ **-r** ] [ *hostname* ] .

---

### ARGUMENTS

**-p** *port* Specify the port number instead of the default SMTP port (number 25) as the next argument.

**-r** Raw mode: disable the default line buffering and input handling. This gives you a similar effect as telnet to port number 25 (not very useful).

### FILES

/usr/lib/sendmail.hf help file for SMTP commands

### SEE ALSO

sendmail

# mkdisknod

## make special files for standard and scsi disk devices

Creates a special file in /dev with the appropriate major number and minor number to permit access to a specified disk partition. Default special files are created for both block and raw (character) mode for partitions b, d, f, h, p, q, r, s, w, x, and z.

---

**/dev/mkdisknod** *disk controller_num disk_num*

---

### ARGUMENTS

*disk* Type of disk:

| | |
|---|---|
| **xdv** | Xylogics VMEbus disk |
| **sd** | SCSI disk |

*controller*
disk controller; an integer in the range 0–15.

*unit*
disk unit; an integer in the range 0–3.

### DESCRIPTION

The system administrator can change the disk partitions created by the csh script mkdisknod. Instructions for modifying are included in the script.

### EXAMPLES

The following example makes special files to access the VMEbus disk on unit 0 of controller 2:

```
# /dev/mkdisknod xdv 2 0
```

The following example makes special files to access the SCSI disk on unit 0 of controller 2:

```
# /dev/mkdisknod sd 2 0
```

### SEE ALSO

mknod

# mkfile

## create NFS swap file

Creates one or more files that are suitable for use as NFS-mounted swap areas. The sticky bit is set and the file is padded with zeroes by default.

---

**mkfile [-nv]** *size*[**k**|**b**|**m**] *filename* ...

---

### ARGUMENTS

**-n** creates an empty file. The size is noted, but disk blocks are not allocated until data is written to them.

**-v** (verbose) reports the names and sizes of created files.

*size*[**k**|**b**|**m**] size of the file in bytes (default), or units of 512 bytes (b), 1000 bytes (k), or 1,048,576 bytes (m).

*filename* name of a file.

### WARNINGS

If a client's swap file is removed and recreated, it must be re-exported before the client can access it. This action may only be done when the client is not running.

### SEE ALSO

exportfs

# mkhosts

## generate hashed host table

Generates the hashed host database used by the library routines gethostbyaddr and gethostbyname. Mkhosts is not used if host name translation is performed by the administration command named.

---

## /etc/mkhosts [-v] *hostfile*

---

### ARGUMENTS

-v (verbose) lists each entry as it is added.

### DESCRIPTION

The file *hostfile* is usually /etc/hosts and in any case must be in the format of /etc/hosts (see the file format for hosts). Mkhost generates database (DBM) files named *hostfile*.pag and *hostfile*.dir. This command leaves the ASCII version of /etc/hosts as it was and creates the DBM version.

The etc/hosts file has an optional new format that allows the system to read it more quickly. The file continues to exist in its regular ASCII format, but now may also reside in a DBM (database) format. If your system has large files for /etc/hosts you may want to use the new format.

When /etc/hosts is read by a system call, the system first looks to see if there is a DBM version of the file. If so, the system reads that version and does not look at the ASCII version. Therefore, whenever you have a DBM version of the file, you must keep it up to date with the ASCII version. Run /etc/mkhosts immediately after editing /etc/hosts.

If there is no DBM format for the file, the system reads the regular ASCII version. (For the sake of compatibility, NIS servers do not read the DBM versions of these files.)

### DIAGNOSTICS

Mkhosts exits with a non-zero exit code if any errors are detected.

### FILES

*hostfile*.pag – database file
*hostfile*.dir – database file
*hostfile*.new.pag – database file
*hostfile*.new.dir – database file

### SEE ALSO

gettable, htable, named
*Concentrix-2800 System Reference Manual*: *System Calls and C Library Routines*: gethostbyname
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: hosts

# mklost+found

## make lost+found directory for fsck

Creates a directory named lost+found in the current directory. Also creates and removes a number of empty files inside the directory so that empty slots exist for /etc/fsck. Typically, /etc/makefs creates a `lost+found` directory in the root of a new file system, and this command is not typed directly.

---

**/etc/mklost+found**

---

### SEE ALSO

fsck, makefs

# mknewdevs

## create default devices for initial installation

Creates default device entries during initial installation.

---

## /dev/mknewdevs

---

The script creates the following devices:

```
crw--w----  1 root     tty       0xc1,       0 Jun 15 11:47 console
crw-rw-rw-  1 root     wheel     0xc3,     0x6 Jul  4  1976 full
crw-rw-rw-  1 root     wheel     0xe0,       0 Jul  4  1976 imb0
crw-rw-rw-  1 root     wheel     0xe0,     0x1 Mar 18 11:33 imb1
crw-rw-rw-  1 root     wheel     0xe0,     0x2 Mar 21 06:06 imb2
crw-rw-rw-  1 root     wheel     0xe0,     0x3 Mar 29 08:22 imb3
crw-rw-rw-  1 root     wheel     0xe0,     0x4 Mar 29 08:22 imb4
crw-rw-rw-  1 root     wheel     0xcf,     0x0 Jul  4  1976 klog
crw-rw-rw-  1 root     wheel     0xc3,     0x1 Jul  4  1976 kmem
crw-rw-rw-  1 root     wheel     0xc3,       0 Jul  4  1976 mem
crw-rw-rw-  1 root     wheel     0xa5,   0x201 Mar 18 09:40 nrschs
lrwxrwxrwx  1 root     wheel                 5 Apr 16 13:34 nrschs02 -> nrschs
crw-rw-rw-  1 root     wheel     0xc3,     0x2 Jun 15 11:20 null
crw-rw-rw-  1 root     wheel     0xc5,       0 Jul  4  1976 ptyp0
crw-rw-rw-  1 root     wheel     0xc5,     0x1 Jun 15 11:58 ptyp1
crw-rw-rw-  1 root     wheel     0xc5,     0x2 Jun 15 11:48 ptyp2
crw-rw-rw-  1 root     wheel     0xc5,     0x3 Jun 15 09:57 ptyp3
crw-rw-rw-  1 root     wheel     0xc5,     0x4 Mar  8 09:29 ptyp4
crw-rw-rw-  1 root     wheel     0xc5,     0x5 Mar  8 09:29 ptyp5
crw-rw-rw-  1 root     wheel     0xc5,     0x6 Mar  8 09:29 ptyp6
crw-rw-rw-  1 root     wheel     0xc5,     0x7 Mar  8 09:29 ptyp7
crw-rw-rw-  1 root     wheel     0xc5,     0x8 Mar  8 09:29 ptyp8
crw-rw-rw-  1 root     wheel     0xc5,     0x9 Mar  8 09:29 ptyp9
crw-rw-rw-  1 root     wheel     0xc5,     0xa Mar  8 09:31 ptypa
crw-rw-rw-  1 root     wheel     0xc5,     0xb Mar  8 09:31 ptypb
crw-rw-rw-  1 root     wheel     0xc5,     0xc Mar  8 09:31 ptypc
crw-rw-rw-  1 root     wheel     0xc5,     0xd Mar  8 09:31 ptypd
crw-rw-rw-  1 root     wheel     0xc5,     0xe Mar  8 09:31 ptype
crw-rw-rw-  1 root     wheel     0xc5,     0xf Mar  8 09:31 ptypf
crw-rw-rw-  1 root     wheel     0xc5,    0x10 Apr  9 15:59 ptyq0
crw-rw-rw-  1 root     wheel     0xc5,    0x11 Apr  9 16:00 ptyq1
crw-rw-rw-  1 root     wheel     0xc5,    0x12 Apr  9 16:00 ptyq2
crw-rw-rw-  1 root     wheel     0xc5,    0x13 Apr  9 16:04 ptyq3
crw-rw-rw-  1 root     wheel     0xc5,    0x14 Apr  9 16:05 ptyq4
crw-rw-rw-  1 root     wheel     0xc5,    0x15 Apr  9 16:10 ptyq5
crw-rw-rw-  1 root     wheel     0xc5,    0x16 Apr  9 16:11 ptyq6
crw-rw-rw-  1 root     wheel     0xc5,    0x17 Apr 10 08:47 ptyq7
crw-rw-rw-  1 root     wheel     0xc5,    0x18 Apr 10 08:47 ptyq8
crw-rw-rw-  1 root     wheel     0xc5,    0x19 Apr 10 08:47 ptyq9
crw-rw-rw-  1 root     wheel     0xc5,    0x1a Apr 10 08:47 ptyqa
crw-rw-rw-  1 root     wheel     0xc5,    0x1b Apr 10 08:47 ptyqb
crw-rw-rw-  1 root     wheel     0xc5,    0x1c Apr 10 08:47 ptyqc
crw-rw-rw-  1 root     wheel     0xc5,    0x1d Apr 10 08:47 ptyqd
crw-rw-rw-  1 root     wheel     0xc5,    0x1e Apr 10 08:47 ptyqe
crw-rw-rw-  1 root     wheel     0xc5,    0x1f Apr 10 08:47 ptyqf
lrwxrwxrwx  1 root     wheel                 5 Apr 16 13:34 rmt0 -> rschs
lrwxrwxrwx  1 root     wheel                 5 Apr 16 13:34 rmt8 -> rschs
lrwxrwxrwx  1 root     wheel                 5 Apr 16 13:34 rmt12 -> nrschs
crw-rw-rw-  1 root     wheel     0xa5,   0x200 Mar 18 09:40 rschs
```

```
lrwxrwxrwx  1 root     wheel                5 Apr 16 13:34 rschs02 -> rschs
crw-r--r--  1 root     wheel     0xc8,      0 Mar 15 17:58 rstripe0
crw-r--r--  1 root     wheel     0xc8,    0x1 Mar 15 17:58 rstripe1
crw-r--r--  1 root     wheel     0xc8,    0x2 Mar 15 17:58 rstripe2
crw-r--r--  1 root     wheel     0xc8,    0x3 Mar 15 17:58 rstripe3
crw-r--r--  1 root     wheel     0xc8,    0x4 Mar 15 17:58 rstripe4
cr--r--r--  1 root     wheel     0x60,    0x1 Mar 18 10:27 rxdv00b
cr--r--r--  1 root     wheel     0x60,    0x3 Jul  4  1976 rxdv00d
cr--r--r--  1 root     wheel     0x60,    0x5 Jul  4  1976 rxdv00f
cr--r--r--  1 root     wheel     0x60,    0xf Mar 18 10:27 rxdv00p
cr--r--r--  1 root     wheel     0x60,   0x10 Mar 18 10:27 rxdv00q
cr--r--r--  1 root     wheel     0x60,   0x11 Jul  4  1976 rxdv00r
cr--r--r--  1 root     wheel     0x60,   0x12 Mar 18 10:27 rxdv00s
cr--r--r--  1 root     wheel     0x60,   0x16 Mar 18 10:27 rxdv00w
cr--r--r--  1 root     wheel     0x60,   0x17 Mar 18 10:27 rxdv00x
cr--r--r--  1 root     wheel     0x60,   0x19 Mar 18 10:27 rxdv00z
cr--r--r--  1 root     wheel     0x60,  0x101 Mar 18 10:28 rxdv01b
cr--r--r--  1 root     wheel     0x60,  0x103 Jul  4  1976 rxdv01d
cr--r--r--  1 root     wheel     0x60,  0x105 Jul  4  1976 rxdv01f
cr--r--r--  1 root     wheel     0x60,  0x107 Mar 18 10:28 rxdv01h
cr--r--r--  1 root     wheel     0x60,  0x10f Mar 18 10:28 rxdv01p
cr--r--r--  1 root     wheel     0x60,  0x110 Mar 18 10:28 rxdv01q
cr--r--r--  1 root     wheel     0x60,  0x111 Jul  4  1976 rxdv01r
cr--r--r--  1 root     wheel     0x60,  0x112 Mar 18 10:28 rxdv01s
cr--r--r--  1 root     wheel     0x60,  0x116 Mar 18 10:28 rxdv01w
cr--r--r--  1 root     wheel     0x60,  0x117 Mar 18 10:28 rxdv01x
cr--r--r--  1 root     wheel     0x60,  0x119 Mar 18 10:28 rxdv01z
cr--r--r--  1 root     wheel     0x60, 0x10001 Mar 18 10:19 rxdv10b
cr--r--r--  1 root     wheel     0x60, 0x10003 Mar 18 10:19 rxdv10d
cr--r--r--  1 root     wheel     0x60, 0x10007 Mar  7 17:35 rxdv10h
cr--r--r--  1 root     wheel     0x60, 0x10016 Mar 18 10:19 rxdv10w
cr--r--r--  1 root     wheel     0x60, 0x10019 Mar 18 10:19 rxdv10z
cr--r--r--  1 root     wheel     0x60, 0x10101 Mar 18 10:19 rxdv11b
cr--r--r--  1 root     wheel     0x60, 0x10103 Mar 18 10:19 rxdv11d
cr--r--r--  1 root     wheel     0x60, 0x10107 Mar 14 15:27 rxdv11h
crw-rw-rw-  1 root     wheel     0x60, 0x10114 Apr 16 08:12 rxdv11u
crw-rw-rw-  1 root     wheel     0x60, 0x10115 Apr 16 08:12 rxdv11v
cr--r--r--  1 root     wheel     0x60, 0x10116 Mar 18 10:19 rxdv11w
cr--r--r--  1 root     wheel     0x60, 0x10119 Mar 18 10:19 rxdv11z
cr--r--r--  1 root     wheel     0x60, 0x20001 Mar 18 10:19 rxdv20b
cr--r--r--  1 root     wheel     0x60, 0x20003 Mar 18 10:19 rxdv20d
cr--r--r--  1 root     wheel     0x60, 0x20007 Mar 14 15:29 rxdv20h
cr--r--r--  1 root     wheel     0x60, 0x20014 Mar 21 21:27 rxdv20u
cr--r--r--  1 root     wheel     0x60, 0x20015 Mar 21 21:28 rxdv20v
cr--r--r--  1 root     wheel     0x60, 0x20016 Mar 18 10:19 rxdv20w
cr--r--r--  1 root     wheel     0x60, 0x20019 Mar 18 10:19 rxdv20z
cr--r--r--  1 root     wheel     0x60, 0x20101 Mar 18 10:19 rxdv21b
cr--r--r--  1 root     wheel     0x60, 0x20103 Mar 18 10:19 rxdv21d
cr--r--r--  1 root     wheel     0x60, 0x20107 Mar 14 15:30 rxdv21h
cr--r--r--  1 root     wheel     0x60, 0x20116 Mar 18 10:19 rxdv21w
cr--r--r--  1 root     wheel     0x60, 0x20119 Mar 18 10:19 rxdv21z
cr--r--r--  1 root     wheel     0x60, 0x30001 Mar 18 10:19 rxdv30b
cr--r--r--  1 root     wheel     0x60, 0x30003 Mar 18 10:19 rxdv30d
cr--r--r--  1 root     wheel     0x60, 0x30007 Mar 14 15:29 rxdv30h
cr--r--r--  1 root     wheel     0x60, 0x30014 Mar 21 21:27 rxdv30u
cr--r--r--  1 root     wheel     0x60, 0x30015 Mar 21 21:28 rxdv30v
cr--r--r--  1 root     wheel     0x60, 0x30016 Mar 18 10:19 rxdv30w
cr--r--r--  1 root     wheel     0x60, 0x30019 Mar 18 10:19 rxdv30z
cr--r--r--  1 root     wheel     0x60, 0x30101 Mar 18 10:19 rxdv31b
cr--r--r--  1 root     wheel     0x60, 0x30103 Mar 18 10:19 rxdv31d
cr--r--r--  1 root     wheel     0x60, 0x30107 Mar 14 15:30 rxdv31h
cr--r--r--  1 root     wheel     0x60, 0x30116 Mar 18 10:19 rxdv31w
cr--r--r--  1 root     wheel     0x60, 0x30119 Mar 18 10:19 rxdv31z
brw-r--r--  1 root     wheel     0xc8,      0 Mar 15 17:58 stripe0
brw-r--r--  1 root     wheel     0xc8,    0x1 Mar 15 17:58 stripe1
brw-r--r--  1 root     wheel     0xc8,    0x2 Mar 15 17:58 stripe2
```

```
brw-r--r--  1 root     wheel    0xc8,     0x3 Mar 15 17:58 stripe3
brw-r--r--  1 root     wheel    0xc8,     0x4 Mar 15 17:58 stripe4
crw-rw-rw-  1 root     wheel    0xc2,       0 Jul  4  1976 tty
crw-rw-rw-  1 root     wheel    0xc1,       0 Jul  4  1976 ttyR0
crw--w--w-  1 root     wheel    0x98,       0 Jun  8 10:44 tty00
crw--w--w-  1 root     wheel    0x98,     0x1 Jun  8 10:44 tty01
crwxrwxrwx  1 root     wheel    0x98,     0x2 Jun  8 10:44 tty02
crwxrwxrwx  1 root     wheel    0x98,     0x3 Jun 15 11:45 tty03
crw--w--w-  1 root     wheel    0x98,     0x4 Jun  8 10:44 tty04
crw--w--w-  1 root     wheel    0x98,     0x5 Jun  8 10:44 tty05
crw--w--w-  1 root     wheel    0x98,     0x6 Jun  8 10:44 tty06
crw--w--w-  1 root     wheel    0x98,     0x7 Jun  8 10:44 tty07
crw--w--w-  1 root     wheel    0x98,     0x8 Jun  8 10:44 tty08
crw--w--w-  1 root     wheel    0x98,     0x9 Jun  8 10:44 tty09
crw--w--w-  1 root     wheel    0x98,     0xa Jun  8 10:44 tty0a
crw--w--w-  1 root     wheel    0x98,     0xb Jun  8 10:44 tty0b
crw--w--w-  1 root     wheel    0x98,     0xc Jun  8 10:44 tty0c
crw--w--w-  1 root     wheel    0x98,     0xd Jun  8 10:44 tty0d
crw--w--w-  1 root     wheel    0x98,     0xe Jun  8 10:44 tty0e
crw--w--w-  1 root     wheel    0x98,     0xf Jun  8 10:44 tty0f
crw--w--w-  1 root     staff    0xc4,       0 Jun 15 11:45 ttyp0
crw--w----  1 perron   tty      0xc4,     0x1 Jun 15 11:57 ttyp1
crw--w--w-  1 lour     hocs     0xc4,     0x2 Jun 15 11:48 ttyp2
crw--w----  1 larsen   tty      0xc4,     0x3 Jun 15 09:57 ttyp3
crw-rw-rw-  1 root     wheel    0xc4,     0x4 Mar  8 09:29 ttyp4
crw-rw-rw-  1 root     wheel    0xc4,     0x5 Mar  8 09:29 ttyp5
crw-rw-rw-  1 garcia   tty      0xc4,     0x6 Mar  8 09:29 ttyp6
crw-rw-rw-  1 crickett tty      0xc4,     0x7 Mar  8 09:29 ttyp7
crw-rw-rw-  1 irene    tty      0xc4,     0x8 Mar  8 09:29 ttyp8
crw-rw-rw-  1 root     wheel    0xc4,     0x9 Mar  8 09:29 ttyp9
crw-rw-rw-  1 root     wheel    0xc4,     0xa Mar  8 09:31 ttypa
crw-rw-rw-  1 root     wheel    0xc4,     0xb Mar  8 09:31 ttypb
crw-rw-rw-  1 root     wheel    0xc4,     0xc Mar  8 09:31 ttypc
crw-rw-rw-  1 lee      tty      0xc4,     0xd Mar  8 09:31 ttypd
crw-rw-rw-  1 root     wheel    0xc4,     0xe Mar  8 09:31 ttype
crw-rw-rw-  1 dolan    tty      0xc4,     0xf Mar  8 09:31 ttypf
crw-rw-rw-  1 root     wheel    0xc4,    0x10 Apr  9 15:58 ttyq0
crw-rw-rw-  1 root     wheel    0xc4,    0x11 Apr  9 16:00 ttyq1
crw-rw-rw-  1 ajg      tty      0xc4,    0x12 Apr  9 16:09 ttyq2
crw-rw-rw-  1 root     wheel    0xc4,    0x13 Apr  9 16:05 ttyq3
crw-rw-rw-  1 root     wheel    0xc4,    0x14 Apr  9 16:05 ttyq4
crw-rw-rw-  1 root     wheel    0xc4,    0x15 Apr  9 16:10 ttyq5
crw-rw-rw-  1 billc    tty      0xc4,    0x16 Apr  9 16:11 ttyq6
crw-rw-rw-  1 mikep    tty      0xc4,    0x17 Apr 10 08:47 ttyq7
crw-rw-rw-  1 root     wheel    0xc4,    0x18 Apr 10 08:47 ttyq8
crw-rw-rw-  1 root     wheel    0xc4,    0x19 Apr 10 08:47 ttyq9
crw-rw-rw-  1 root     wheel    0xc4,    0x1a Apr 10 08:47 ttyqa
crw-rw-rw-  1 root     wheel    0xc4,    0x1b Apr 10 08:47 ttyqb
crw-rw-rw-  1 root     wheel    0xc4,    0x1c Apr 10 08:47 ttyqc
crw-rw-rw-  1 root     wheel    0xc4,    0x1d Apr 10 08:47 ttyqd
crw-rw-rw-  1 root     wheel    0xc4,    0x1e Apr 10 08:47 ttyqe
crw-rw-rw-  1 root     wheel    0xc4,    0x1f Apr 10 08:47 ttyqf
br--r--r--  1 root     wheel    0x60,     0x1 Mar 18 10:27 xdv00b
br--r--r--  1 root     wheel    0x60,     0x3 Jul  4  1976 xdv00d
br--r--r--  1 root     wheel    0x60,     0x5 Jul  4  1976 xdv00f
br--r--r--  1 root     wheel    0x60,     0xf Mar 18 10:27 xdv00p
br--r--r--  1 root     wheel    0x60,    0x10 Mar 18 10:27 xdv00q
br--r--r--  1 root     wheel    0x60,    0x11 Jul  4  1976 xdv00r
br--r--r--  1 root     wheel    0x60,    0x12 Mar 18 10:27 xdv00s
br--r--r--  1 root     wheel    0x60,    0x16 Mar 18 10:27 xdv00w
br--r--r--  1 root     wheel    0x60,    0x17 Mar 18 10:27 xdv00x
br--r--r--  1 root     wheel    0x60,    0x19 Mar 18 10:27 xdv00z
br--r--r--  1 root     wheel    0x60,   0x101 Mar 18 10:28 xdv01b
br--r--r--  1 root     wheel    0x60,   0x103 Jul  4  1976 xdv01d
br--r--r--  1 root     wheel    0x60,   0x105 Jul  4  1976 xdv01f
br--r--r--  1 root     wheel    0x60,   0x107 Mar 18 10:28 xdv01h
```

```
br--r--r--  1 root     wheel    0x60,   0x10f Mar 18 10:28 xdv01p
br--r--r--  1 root     wheel    0x60,   0x110 Mar 18 10:28 xdv01q
br--r--r--  1 root     wheel    0x60,   0x111 Jul  4  1976 xdv01r
br--r--r--  1 root     wheel    0x60,   0x112 Mar 18 10:28 xdv01s
br--r--r--  1 root     wheel    0x60,   0x116 Mar 18 10:28 xdv01w
br--r--r--  1 root     wheel    0x60,   0x117 Mar 18 10:28 xdv01x
br--r--r--  1 root     wheel    0x60,   0x119 Mar 18 10:28 xdv01z
br--r--r--  1 root     wheel    0x60, 0x10001 Mar 18 10:19 xdv10b
br--r--r--  1 root     wheel    0x60, 0x10003 Mar 18 10:19 xdv10d
br--r--r--  1 root     wheel    0x60, 0x10007 Mar  7 17:35 xdv10h
br--r--r--  1 root     wheel    0x60, 0x10016 Mar 18 10:19 xdv10w
br--r--r--  1 root     wheel    0x60, 0x10019 Mar 18 10:19 xdv10z
br--r--r--  1 root     wheel    0x60, 0x10101 Mar 18 10:19 xdv11b
br--r--r--  1 root     wheel    0x60, 0x10103 Mar 18 10:19 xdv11d
br--r--r--  1 root     wheel    0x60, 0x10107 Mar 14 15:27 xdv11h
brw-rw-rw-  1 root     wheel    0x60, 0x10114 Apr 16 08:13 xdv11u
brw-rw-rw-  1 root     wheel    0x60, 0x10115 Apr 16 08:13 xdv11v
br--r--r--  1 root     wheel    0x60, 0x10116 Mar 18 10:19 xdv11w
br--r--r--  1 root     wheel    0x60, 0x10119 Mar 18 10:19 xdv11z
br--r--r--  1 root     wheel    0x60, 0x20001 Mar 18 10:19 xdv20b
br--r--r--  1 root     wheel    0x60, 0x20003 Mar 18 10:19 xdv20d
br--r--r--  1 root     wheel    0x60, 0x20007 Mar 14 15:29 xdv20h
br--r--r--  1 root     wheel    0x60, 0x20014 Mar 21 21:27 xdv20u
br--r--r--  1 root     wheel    0x60, 0x20015 Mar 21 21:28 xdv20v
br--r--r--  1 root     wheel    0x60, 0x20016 Mar 18 10:19 xdv20w
br--r--r--  1 root     wheel    0x60, 0x20019 Mar 18 10:19 xdv20z
br--r--r--  1 root     wheel    0x60, 0x20101 Mar 18 10:19 xdv21b
br--r--r--  1 root     wheel    0x60, 0x20103 Mar 18 10:19 xdv21d
br--r--r--  1 root     wheel    0x60, 0x20107 Mar 14 15:31 xdv21h
br--r--r--  1 root     wheel    0x60, 0x20116 Mar 18 10:19 xdv21w
br--r--r--  1 root     wheel    0x60, 0x20119 Mar 18 10:19 xdv21z
```

## SEE ALSO

mkdisknod, mknod, mkptynod, mkttynod, mktape

# mknod

## build special file or make a named pipe

Makes a special file in /dev for accessing a device or makes a named pipe.

---

**/etc/mknod** *name* [c] [b] *major minor*

**/etc/mknod** *name* **p**

---

### ARGUMENTS

*name* specifies the name of the device. The name is arbitrary but certain conventions are usually followed.

**b** specifies the access mode as block (cooked; file oriented access).

**c** specifies the access mode as character (raw; character-by-character access).

**p** specifies a named pipe (fifo).

*major* is a number that specifies the device type and controller number. The major number must be appropriate for the device being accessed. The number can be specified in decimal or hexadecimal; specify hexadecimal with the prefix 0x.

*minor* is a number that specifies the unit number where necessary and other device specific information such as tape density. The number can be specified in decimal or hexadecimal; specify hexadecimal with the prefix 0x.

### EXAMPLES

The following example makes a special file named /dev/rxtv00h for a VMEbus magnetic tape drive, on controller 0 drive 0.

```
# /etc/mknod /dev/rxtv00h c 0x80 0
# chmod 666 /dev/rxtv00h
```

The following example creates a named pipe (fifo) *pipe* through which two unrelated processes can subsequently communicate:

```
# /etc/mknod pipe p
```

### SEE ALSO

Section 3.1
mkdisknod, mktape, mkttynod
*Concentrix-2800 System Reference Manual: System Calls and C Library Routines*: mknod

# mkpasswd

## generate hashed password table

Generates the hashed password data base used by the library routines getpwnam and getpwuid.

---

**/etc/mkpasswd [-v]** *passwdfile*

---

### ARGUMENTS

-v (verbose) lists each entry as it is added.

### DESCRIPTION

The file *passwdfile* must be in the format of /etc/passwd. Mkpasswd generates database (DBM) files named *passwdfile*.pag and *passwdfile*.dir.

When /etc/passwd is read by a system call, the system first looks to see if there is a DBM version of the file. If so, the system reads that version. If there is not a DBM format for the file, the system reads the regular ASCII version. (For the sake of compatibility, NIS servers do not read the DBM versions of the file.)

When the system finds the DBM version of the files, it never looks in the ASCII version. Therefore, whenever you have a DBM version of these files (required for /etc/passwd), you must keep it up to date with the ASCII version. Use vipw or passwd to modify /etc/passwd.

### DIAGNOSTICS

Mkpasswd exits with a non-zero exit code if any errors are detected.

### FILES

| | |
|---|---|
| passwdfile | source file in the format of /etc/passwd |
| passwdfile.pag | data base file |
| passwdfile.dir | data base file |

### SEE ALSO

vipw
*Concentrix-2800 System Reference Manual: System Calls and C Library Routines*: getpwent
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats:* passwd

# mkproto

## construct a prototype file system

Mkproto bootstraps a new file system. First a new file system is created with makefs. Mkproto then copies files from the old file system into the new file system according to the directions found in the prototype file `proto`.

---

## /usr/etc/mkproto *special proto*

---

### ARGUMENTS

*special* is the name of a special file as described below.

*proto* is the name of a prototype file as described below.

### DESCRIPTION

The prototype file `proto` contains tokens separated by spaces or new lines. The first tokens comprise the specification for the root directory. File specifications consist of tokens giving the mode, the user-id, the group id, and the initial contents of the file. The syntax of the contents field depends on the mode.

The mode token for a file is a six-character string. The first character specifies the type of the file. (The characters **-bcd** specify regular, block special, character special and directory files respectively.) The second character of the type is either **u** or **-** to specify set-user-id mode or not. The third is **g** or **-** for the set-group-id mode. The rest of the mode is a three digit octal number giving the owner, group, and other read, write, execute permissions, see *chmod*.

Two decimal number tokens come after the mode; they specify the user and group ID's of the owner of the file.

If the file is a regular file, the next token is a pathname from which the contents and size are copied.

If the file is a block or character special file, two decimal number tokens follow that give the major and minor device numbers.

If the file is a directory, mkproto makes the entries . and .. and then reads a list of names and (recursively) file specifications for the entries in the directory. The scan is terminated with the token $.

A sample prototype specification follows:

```
d--777 3 1
usr         d--777 3 1
            sh         ---755 3 1 /bin/sh
            ken        d --755 6 1
                       $
            b0         b--644 3 1 0 0
            c0         c--644 3 1 0 0
            $
$
```

## NOTES

- There should be some way to specify links.
- There should be some way to specify bad blocks.
- Mkproto can only be run on virgin file systems. It should be possible to copy files into existent file systems.

## SEE ALSO

fsck, makefs
*Concentrix Commands and Applications Manual*: chmod
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: fs, dir

# mkptynod

## make special files for pseudo terminals

Creates special files in /dev with the appropriate major numbers and minor numbers to provide connections for remote logins. Master pseudo terminal names (major number 0xc5 – 197) are ptyp0 through ptypf, ptyq0 through ptyqf, and so on. Slave pseudo terminal names (major number 0xc4 –196) are ttyp0 through ttypf, ttyq0 through ttyqf, and so on.

The special files are created in your current working directory.

---

## /dev/mkptynod [#] *num*

---

### ARGUMENTS

*#* qualifies *num* as described below.

*num* is an integer in the range 0 –7 that specifies:

# not specified: the number of pseudo terminal groups to create starting with ptyp/ttyp.

# specified: the pseudo terminal group to create, where ptyp/ttyp is group 0.

### EXAMPLES

The following example makes pseudo terminals for ptyp/ttyp, ptyq/ttyq, and ptyr/ttyr:

```
# /dev/mkptynod 3
```

The following example makes pseudo terminals for ptyr/ttyr:

```
# /dev/mkptynod # 3
```

### NOTES

- The number of pseudo terminal ports available in the system is limited by the kernel variable nptys, which can be modified by the command ptynum. The default value is 64, thus limiting the useful range of mkptynod to 0–3. Increases in the nptys variable should be in increments of 16. The mkptynod command can create device entries for 128 pseudo terminals.

### SEE ALSO

mknod, ptynum, mkttynod

# mkstripe

## make striped partition

Assembles disk partitions into a disk stripe.

---

**/etc/mkstripe** *raw-stripe raw-disk raw-disk ...*

**/etc/mkstripe –a**

---

### ARGUMENTS

*raw-stripe* is the raw device name of the disk stripe.

*raw-disk* is the raw device name of a raw disk partition. Two or more partitions must be named.

**–a** reads the file /etc/stripetab and assembles the striped partitions described there.

### EXAMPLES

The following example makes a disk stripe from the partitions /dev/rxdv00f and /dev/rxdv10f:

```
# /etc/mkstripe /dev/rstripe0 /dev/rxdv00f /dev/rxdv10f
```

The following example makes disk stripes for striped devices in /etc/stripetab:

```
# /etc/mkstripe -a
```

### FILES

/etc/stripetab striped partition database

Each line of /etc/stripetab defines one disk stripe. Specify the raw device name of the striped device followed by the raw device names of the partitions that will make up the striped file system. The names must be separated by spaces or tabs. A comment can be placed in the file by starting a line with a pound sign (#).

### SEE ALSO

makefs, unmkstripe, stripetab

# mkswap

## make node for swapping partition

Makes a node for the swapping partition. The node is used by other Concentrix-2800 utilities, such as *ps* and w, to collect data on processes that are swapped out. Typically, /etc/rc executes mkswap when the system boots.

---

## /bin/mkswap

---

# mktape

## make special files for tapes

Creates eight special files in /dev with the appropriate major number and minor number to permit access to the given controller and unit (rewind and no rewind at each of four densities).

The file creation mask, if set, determines the permissions of the devices that are made. You may want to change the permissions, for example, to 666 so that everyone can access them.

---

**/usr/alliant/mktape** *controller unit brand*

---

### ARGUMENTS

*controller* an integer in the range 0–11.

*unit* an integer in the range 0–7.

*brand* one of the following (as supplied by Alliant):

0 = telex50
1 = cipher
2 = fujitsu
3 = telex125
4 = keystone
5 = kennedy9100
6 = pertec
7 = anritsu
9 = stc
a = thorn
b = special1
c = special2
d = special3

### EXAMPLES

The following example makes the special files /dev/rxtv10h, /dev/nrxtv10h, /dev/rxtv10m, /dev/nrxtv10m, /dev/rxtv10l, /dev/nrxtv10l, /dev/rxtv10x, and /dev/nrxtv10x, and associates the devices with a cipher tape drive on unit 0 of controller 1:

```
# /dev/mktape 1 0 cipher
# chmod 666 /dev/*rxtv10*
```

### SEE ALSO

Section 3.1, mknod

# mkttynod

## make special files for terminals

Creates special files in /dev with the appropriate major numbers and minor numbers to permit access to terminal ports connected to a specified multiplexor.

---

## /dev/mkttynod xmv *controller type*

---

### ARGUMENTS

| | |
|---|---|
| xmv | VMEbus multiplexor |
| *controller* | multiplexor; an integer in the range 0–7. |
| *type* | one of two values: |
| | **t** terminal (no modem) or outgoing modem |
| | **m** incoming modem |

### EXAMPLES

The following example makes 16 terminal ports for incoming modems in the range tty30 to tty3f:

```
# /dev/mkttynod xmv 3 m
```

### SEE ALSO

mknod, mkptynod

# mount

## mount file systems

Attaches a file system to the file hierarchy at a specified directory. If the specified directory has any contents prior to the mount operation, these remain hidden until the filesystem is once again unmounted. If the filesystem is of the form host:pathname, it is assumed to be an NFS filesystem (type nfs).

---

**/etc/mount**

**/etc/mount -p**

**/etc/mount -a[f][n][v] [-t** *type***]**

**/etc/mount -[f][n][r][v] [-t** *type***] [-o** *options***]** *fsname dir*

**/etc/mount -[f][n][v] [-o** *options***]** *fsname* | *dir*

---

### ARGUMENTS

If invoked with no arguments, mount displays /etc/mtab. If invoked with either a filesystem or directory only, mount searches the file /etc/fstab for a matching entry, and mounts the filesystem indicated in that entry on the indicated directory.

> **NOTE:** Use the -v switch to get error messages. Previously, error messages were reported automatically.

**-a** mounts all the file systems described in /etc/fstab or all the file systems of the specified type. *Fsname* and *dir* are not specified.

**-f** (fake) makes a new /etc/mtab entry, but does not actually mount any file systems.

**-n** Mount the filesystem without making an entry in /etc/mtab.

**-o** specifies mount options, separated by commas:

- **ro** read-only.
- **rw** read-write (default). NFS filesystems that are mounted rw should use the **hard** option.
- **suid** set-uid execution permitted (default).
- **nosuid** set-uid execution not permitted.
- **grpid** Create files with BSD semantics for the propagation of the group ID. Under this option, files inherit the GID of the directory in which they are created, regardless of the directory's set-GID bit.
- **noauto** automatic mounting (-a) not permitted.

**remount** If the file system is currently mounted, and if the entry in /etc/fstab specifies that it is to be mounted read-write or **rw** was specified along with **remount**, remount the file system making it read-write. If the entry in /etc/fstab specifies that it is to be mounted read-only and **rw** was not specified, the file system is not remounted. If the file system is not currently mounted, an error results.

**quota|noquota**
(ufs/tfs file systems only) Usage limits are enforced, or not enforced. The default is **noquota.**

**bg** (nfs file systems only) if the first mount attempt fails, retry in the background.

**fg** (nfs file systems only) if the first mount attempt fails, retry in the foreground (default).

**retry=*n*** (nfs file systems only) number of times to retry (defaults to 1).

**rsize=*n*** (nfs file systems only) size of read buffer in bytes (defaults to 8192).

**wsize=*n*** (nfs file systems only) size of write buffer in bytes (defaults to 8192).

**timeo=*n*** (nfs file systems only) NFS time-out in tenths of a second (defaults to 7). On a mounted file system, NFS waits the indicated amount of time for a response to a request. If no response occurs, NFS multiplies *n* by 2 and retries.

**retrans=*n*** (nfs file systems only) number of NFS retransmissions (defaults to 3). On a mounted file system, if a request is sent and no response occurs, NFS multiplies the time-out time by 2 and retries the transmission the specified number of times.

**port=*n*** (nfs file systems only) server IP port number (defaults to NFS_PORT).

**hard** (nfs file systems only) causes mount to try forever if the mountd server does not respond (default). On a hard mounted file system, if a request fails (the time-out and retransmission limits are exceeded), an error message is issued and the request is tried again. File systems mounted rw should use the hard option.

**soft** (nfs file systems only) causes mount to fork a background copy to retry forever if the mountd server does not respond. On a soft mounted file system, if a request fails (the time-out and retransmission limits are exceeded), the access request fails with ETIMEDOUT.

**intr** (nfs file systems only) permits keyboard interrupts on hard mounts, so that a process waiting for a response on a hard mounted file system can be killed.

**secure** (nfs only) Use a more secure protocol for NFS transactions. The secure option must be given if the server requires secure mounting for the file system.

**acregmin=n**
Hold cached attributes for at least n seconds after file modification. (Defaults to 3.)

**acregmax=n**
Hold cached attributes for no more than n seconds after file modification. (Defaults to 60.)

**acdirmin=n**
Hold cached attributes for at least n seconds after directory update. (Defaults to 30.)

**acdirmax=n**
Hold cached attributes for no more than n seconds after directory update. (Defaults to 60.)

**actimeo=n**
Set *min* and *max* times for regular files and directories to n seconds.

The attribute cache retains file attributes on the client. Attributes for a file are assigned a time to be flushed. If the file is modified before the flush time, then the flush time is extended by the time since the last modification (under the assumption that files that changed recently are likely to change soon). There is a minimum and maximum flush time extension for regular files and for directories. Setting **actimeo=n** extends flush time by n seconds for both regular files and directories.

**-p** prints the list of mounted file systems in a format suitable for use in /etc/fstab.

**-r** mounts the specified file system read-only. This is a shorthand for **-o ro**.

Physically write-protected and magnetic tape file systems must be mounted read-only. Otherwise, errors occur when access times are updated, whether or not any explicit write is attempted.

**-t** specifies the file system type. This option is effective only with **-a** to qualify the file systems being mounted. Named local file systems are mounted as the types that they are; named remote files systems are mounted as **nfs**.

**4.3** specifies a 4.3BSD file system. (4.2 is treated the same as 4.3)

**tfs** specifies a track-allocated file system.

**nfs** specifies a remote file system.

**-v** (verbose) displays a message naming the file system being mounted and the mount point.

*fsname* (file system name) specifies a name of one of the following forms:

**/dev/***part* block device name of a disk partition on the local system. The partition must contain a file system (see **makefs**). The file system type is **4.3** or **tfs** depending on the contents of the partition.

*host:path* mount point (*host*) of a file system (*path*) on a remote system. The file system type is **nfs**.

*dir* (mount point) specifies an existing directory in the root file system that becomes the root of the newly mounted file system. Any existing contents of *dir* are hidden until the file system is unmounted.

If *fsname* or *dir* is specified alone, an entry containing that name must exist in /etc/fstab, and the missing name is derived from the entry.

## EXAMPLES

The following example mounts all file systems named in /etc/fstab and prints messages:

```
# /etc/mount -a -v
```

The following example mounts all file systems of type tfs named in /etc/fstab and prints messages:

```
# /etc/mount -a -t -v tfs
```

The following example mounts the file system in the local disk partition /dev/xdv01h at the mount point /projectx and prints messages:

```
# /etc/mount -v /dev/xdv01h /projectx
```

The following example mounts a file system at the mount point /projectx and prints messages. An entry naming /projectx as a mount point must be present in /etc/fstab.

```
# /etc/mount -v /projectx
```

The following example mounts a file system from system sysx at the mount point /projectx and prints messages:

```
# /etc/mount -v sysx:/projectx /projectx
```

The following example saves /etc/fstab and writes a new file system table based on the current mounted file systems and prints messages:

```
# cp /etc/fstab /etc/fstab.old
# /etc/mount -p -v > /etc/fstab
```

## FILES

| | |
|---|---|
| /etc/mtab | mounted file system table |
| /etc/fstab | static information about file systems |

## NOTES

- Mounting a corrupted file system crashes the system.
- 4.2 filesystem types are silently considered as 4.3BSD.

## SEE ALSO

makefs, mountd, nfsd, umount
*Concentrix Commands and Applications Manual*, df
*Concentrix-2800 System Reference Manual: System Calls and C Library Routines*: mkdir, mount, nfsmount, umount, open
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: fstab, mtab

# mountd

## NFS mount request server

Responds to file system mount requests from other systems. Reads the file /etc/xtab to determine which file systems are available to which machines. Provides information as to which client systems have file systems mounted. This information can be printed using the showmount command.

The mountd daemon is started as needed by /etc/inetd.

---

## /usr/etc/rpc.mountd [-n]

---

### ARGUMENTS

**-n** does not ensure that client users are root. Permits older versions of NFS to work at the sacrifice of some security.

### FILES

| | |
|---|---|
| /etc/xtab | mounted external file system table |

### SEE ALSO

inetd, showmount
*Concentrix-2800 System Reference Manual*: *Miscellaneous Routines, Special Files, and File Formats*: exports, services

# named

## Internet domain name server

Named is the Internet domain name server as specified in RFC883.

The named daemon is typically invoked by /etc/rc.local. (As distributed, the named invocation is commented out in /etc/rc.local. Uncomment the invocation if you want named to start automatically with the system.)

---

**named** [**-d** *debuglevel*] [**-p** *port#*] [{**-b**} *bootfile*]

---

### ARGUMENTS

**-d** *debuglevel*
prints debugging information at the specified message level.

**-p** *port#* uses the specified port number. The default is the standard port number as listed in /etc/services.

**-b** Use an alternate boot file. This is optional and allows you to specify a file with a leading dash.

*bootfile* specifies the boot file; defaults to /etc/named.boot.

### DESCRIPTION

The boot file contains information about where the name server gets its initial data. If multiple boot files are specified, only the last is used. Lines in the boot file cannot be continued on subsequent lines. The following is a small example:

```
;
;       boot file for name server
;
; type          domain          source host/            backup file

cache       .                                             root.cache
primary     Berkeley.EDU          berkeley.edu.zone
primary     32.128.IN-ADDR.ARPA   ucbhosts.rev
secondary   CC.Berkeley.EDU       128.32.137.8 128.32.137.3 cc.zone.bak
secondary   6.32.128.IN-ADDR.ARPA 128.32.137.8 128.32.137.3 cc.rev.bak
primary     0.0.127.IN-ADDR.ARPA                          localhost.rev
forwarders 10.0.0.78 10.2.0.78
; slave
```

The "directory" line causes the server to change its working directory to the directory specified. This can be important for the correct processing of $INCLUDE files in primary zone files.

The "cache" line specifies that data in root.cache is to be placed in the backup cache. Its main use is to specify data such as locations of root domain servers. This cache is not

used during normal operation, but is used as "hints" to find the current root servers. The file root.cache is in the same format as berkeley.edu.zone. There can be more than one cache file specified. The cache files are processed in such a way as to preserve the time-to-live's of data dumped out. Data for the root nameservers are kept artificially valid if necessary.

The first "primary" line states that the file berkeley.edu.zone contains authoritative data for the Berkeley.EDU zone. The file berkeley.edu.zone contains data in the master file format described in RFC883. All domain names are relative to the origin, in this case,"Berkeley.EDU" (see below for a more detailed description). The second "primary" line states that the file ucbhosts.rev contains authoritative data for the domain "32.128.IN-ADDR.ARPA," which is used to translate addresses in network 128.32 to hostnames. Each master file should begin with an SOA record for the zone (see below).

The first "secondary" line specifies that all authoritative data under "CC.Berkeley.EDU" is to be transferred from the name server at 128.32.137.8. If the transfer fails it will try 128.32.137.3 and continue trying the addresses, up to 10, listed on this line. The secondary copy is also authoritative for the specified domain. The first non-dotted-quad address on this line will be taken as a filename in which to backup the transfered zone. The name server will load the zone from this backup file if it exists when it boots, providing a complete copy even if the master servers are unreachable. Whenever a new copy of the domain is received by automatic zone transfer from one of the master servers, this file will be updated. The second "secondary" line states that the address-to-hostname mapping for the subnet 128.32.136 should be obtained from the same list of master servers as the previous zone.

The "forwarders" line specifies the addresses of sitewide servers that will accept recursive queries from other servers. If the boot file specifies one or more forwarders, then the server will send all queries for data not in the cache to the forwarders first. Each forwarder will be asked in turn until an answer is returned or the list is exhausted. If no answer is forthcoming from a forwarder, the server will continue as it would have without the forwarders line unless it is in "slave" mode. The forwarding facility is useful to cause a large sitewide cache to be generated on a master, and to reduce traffic over links to outside servers. It can also be used to allow servers to run that do not have access directly to the Internet, but wish to act as though they do.

The "slave" line (shown commented out) is used to put the server in slave mode. In this mode, the server will only make queries to forwarders. This option is normally used on machine that wish to run a server but for physical or administrative reasons cannot be given access to the Internet, but have access to a host that does have access.

The "sortlist" line can be used to indicate networks that are to be preferred over other, unlisted networks. Queries for host addresses from hosts on the same network as the server will receive responses with local network addresses listed first, then addresses on the sort list, then other addresses. This line is only acted on at initial startup. When reloading the nameserver with a SIGHUP, this line will be ignored.

The master file consists of control information and a list of resource records for objects in the zone of the forms:

```
$INCLUDE <filename> <opt_domain>
$ORIGIN <domain>
<domain> <opt_ttl> <opt_class> <type> <resource_record_data>
```

Where *domain* is a period (.) for root, an at sign (@) for the current origin, or a standard domain name. If *domain* is a standard domain name that does not end with a period, the current origin is appended to the domain. Domain names ending with a period are unmodified. The *opt_domain* field is is used to define an origin for the data in an included file. It is equivalent to placing a $ORIGIN statement before the first line of the included file. The field is optional. Neither the opt_domain field nor $ORIGIN statements in the included file modify the current origin for this file. The *opt_ttl field* is an optional integer number for the time-to-live field and defaults to zero. The *opt_class* field is the object address type; currently only one type is supported, IN, for objects connected to the DARPA Internet. The *type* field is one of the following tokens; the data expected in the *resource_record_data* field is in parentheses.

| | |
|---|---|
| A | host address (dotted quad) |
| NS | authoritative name server (domain) |
| MX | mail exchanger (domain) |
| CNAME | canonical name for an alias (domain) |
| SOA | start of a zone of authority (domain of originating host, domain address of maintainer, a serial number and the following parameters in seconds: refresh, retry, expire and minimum TTL (see RFC883)) |
| MB | mailbox domain name (domain) |
| MG | mail group member (domain) |
| MR | mail rename domain name (domain) |
| NULL | null resource record (no format or data) |
| WKS | well know service description (not implemented yet) |
| PTR | domain name pointer (domain) |
| HINFO | host information (cpu_type OS_type) |
| MINFO | mailbox or mail list information (request_domain error_domain) |

Resource records normally end at the end of a line, but may be continued across lines between opening and closing parentheses. Comments are introduced by semicolons and continue to the end of the line.

Each master zone file should begin with an SOA record for the zone. An example SOA record is as follows:

```
@       IN      SOA     ucbvax.Berkeley.EDU. rwh.ucbvax.Berkeley.EDU. (
                                2.89    ; serial
                                10800   ; refresh
                                3600    ; retry
                                3600000 ; expire
                                86400 ) ; minimum
```

The SOA lists a serial number, which should be changed each time the master file is changed. Secondary servers check the serial number at intervals specified by the refresh time in seconds; if the serial number changes, a zone transfer will be done to load the new data. If a master server cannot be contacted when a refresh is due, the retry time specifies the interval at which refreshes should be attempted until successful. If a master server cannot be contacted within the interval given by the expire time, all data from the zone is discarded by secondary servers. The minimum value is the time-to-live used by records in the file with no explicit time-to-live value.

> **NOTE:** The boot file directives "domain" and "suffixes" have been obsoleted by a more useful resolver-based implementation of suffixing for partially qualified domain names. The prior mechanisms could fail under a number of situations, especially when then local nameserver did not have complete information.

## NOTES

The following signals have the specified effect when sent to the server process using the kill command.

SIGHUP causes server to read named.boot and reload database.

SIGINT Dumps current data base and cache to /usr/tmp/named_dump.db.

SIGIOT Dumps statistics data into /usr/tmp/named.stats if the server is compiled -DSTATS. Statistics data is appended to the file.

SIGSYS Dumps the profiling data in /usr/tmp if the server is compiled with profiling (server forks, chdirs and exits).

SIGTERM Dumps the primary and secondary database files. Used to save modified data on shutdown if the server is compiled with dynamic updating enabled.

SIGUSR1 Turns on debugging; each SIGUSR1 increments debug level.

SIGUSR2 Turns off debugging completely.

## FILES

| | |
|---|---|
| /etc/named.boot | name server configuration boot file |
| /etc/named.pid | the process id |
| /usr/tmp/named_dump.db | dump of the name server database |
| /usr/tmp/named.run | debug output |
| /usr/tmp/named.stats | nameserver statistics data |

## SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: kill
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: gethostbyname, signal, res_mkquery, resolver
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: resolver

# ndbootd

## ND boot block server

Sends boot blocks to diskless Sun-2 system clients that request them using the (now obsolete) ND protocol.

---

**/usr/etc/install/ndbootd [-d] [-v]**

---

This server uses the boot block contained in the file /tftpboot/sun2.bb. A client must appear in the ethers and hosts data bases for the request to be served. In determining whether to serve the client, ndbootd checks the /tftpboot directory for a file whose name is the client's IP address in hexadecimal notation. For example, if either /tftpboot/C00901AD or /tftpboot/C00901AD.SUN2 exists, the machine at IP address 192.9.1.173 can be served. This file normally contains the boot program that is sent to the client by tftpd.

Only root can invoke ndbootd.

### ARGUMENTS

**-d** (debug) Display information about ignored packets, retransmissions, and address translation.

**-v** (verbose) show a detailed listing of packets sent and received.

If either debug or verbose is used, all output is sent to the invoking terminal. Otherwise, error output (if any) appears on the console.

### FILES

| | |
|---|---|
| /tftpboot | bootfiles directory |
| /tftpboot/sun2.bb | boot blocks |
| /tftpboot/???????? | boot programs for clients |

### SEE ALSO

tftpd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ethers, hosts

# netdae

## batch network server daemon

Listens for incoming requests for the batch system and invokes the batch server process to service the request. Normally the server is /usr/lib/batch/netdae. Note: this program is invoked by the batch system and is not intended for interactive use.

---

**/usr/lib/batch/netdae [-d] [-n *connections*] [-p *port*] *server* [*arg* ...]**

---

### ARGUMENTS

**-d** causes netdae to print debugging information on the standard error output.

**-n** *connections*
controls the maximum number of simultaneous connections that the daemon allows to exist. Defaults to four.

**-p** *port* specifies the TCP port that netdae waits on.

*server* [*arg* ...]
calls *server* with any optional arguments and -h remotehost. File descriptors 0 and 1 of the server are connected to the network socket, and file descriptor 2 remains unchanged.

### SEE ALSO

batchd, netrecv, netsend
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: qconf

# netrecv

## batch network receive daemon

Processes batch requests from remote systems. A request is accepted on the standard input and a reply code is returned on the standard output. Note: this program is invoked by the batch system and is not intended for interactive use.

---

**/usr/lib/batch/netrecv [-d] [-h *remotehost*]**

---

### ARGUMENTS

**-d** causes netrecv to print debugging information on the standard error output.

**-h** *remotehost*
specifies the host that originated the request. This specification is taken to be more authoritative than the site information in the network protocol. The host information is used to check access permissions (see below).

### DESCRIPTION

If the configuration file /etc/qconf contains the parameter accessfile, the file specified by accessfile is checked to determine if *remotehost* has permission to submit requests to the queue. The format of the access file is one line per allowed host. Each line consists of a host name optionally followed by queue names; the names are separated by white space. If no queues are listed, the request can be submitted to any queue. If queues are listed, the request can only be submitted to one of those queues. If an attempt is made to submit to an invalid queue or an unauthorized queue, the request is rejected.

### FILES

| | |
|---|---|
| /etc/qconf | batch system configuration file |

### SEE ALSO

batchdae, netdae, netsend
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: qconf

# netsend

## batch network send daemon

Processes batch requests for remote systems. A request is accepted on the standard input and a reply code is returned on the standard output. Note: this program is invoked by the batch system and is not intended for interactive use.

---

**/usr/lib/batch/netsend** *remotehost* [**-p** *port*] *remotequeue*

---

### ARGUMENTS

*remotehost*
specifies the host receiving the request. This specification is taken to be more authoritative than the site information in the network protocol.

**-p** *port* specifies the TCP port that netdae waits on.

*remotequeue*
specifies a queue on the remote host.

### FILES

/etc/qconf batch system configuration file

### SEE ALSO

batchdae, netdae, netrecv
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: qconf

# newaliases

## rebuild data base for mail aliases file

(Identical to sendmail -bi.) Rebuilds the random access data base for the binary mail aliases files /usr/lib/aliases.dir and /usr/lib/aliases.pag.

---

## /usr/ucb/newaliases

---

### DESCRIPTION

Run newaliases after changing /usr/lib/aliases to make the changes effective.

### FILES

| | |
|---|---|
| /usr/lib/aliases | aliases file for sendmail |
| /usr/lib/aliases.dir | binary representation of aliases file |
| /usr/lib/aliases.pag | binary representation of aliases file |

### SEE ALSO

sendmail
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: aliases

# newkey

## create a new key in the publickey database

Newkey prompts for the login password of the given *username* and then creates a new public/secret key pair in /etc/publickey encrypted with the login password of the given user.

---

**newkey** [**-h** *hostname* ] [ **-u** *username* ]

---

### ARGUMENTS

**-h** *hostname*
Create a new public key for the superuser at the given *hostname*. Prompts for the root password of the given *hostname*.

**-u** *username*
Create a new public key for the given *username*. Prompts for the NIS password of the given username.

### DESCRIPTION

newkey is normally run by the network administrator on the NIS master machine in order to establish public keys for users and superusers on the network. These keys are needed for using secure RPC or secure NFS.

Use of this program is not required: users may create their own keys using chkey.

### SEE ALSO

keyserv
*Concentrix-2800 Commands and Applications:* chkey, keylogin
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: publickey

# newfs

## construct a new file system (unsupported)

Use makefs.

# nfsd

## NFS daemon

Starts the NFS server daemons that handle client filesystem requests.

Four nfsd daemons are typically invoked by /etc/rc.local. (As distributed, the nfsd invocation is commented out in /etc/rc.local. Uncomment the invocation if you want nfsd to start automatically with the system.)

---

## /etc/nfsd [*nservers*]

---

### ARGUMENTS

*nservers* is the number of file system request daemons to start. The number should be based on the load expected on this server; four is a good number. Defaults to one.

### NOTES

When a file opened by a client is unlinked (by the server), a file with a name of the form .nfs*nnn* is created by the client. When the open file is closed, the .nfs*nnn* is removed. If the client crashes before the file can be closed, the .nfs*nnn* remains.

### SEE ALSO

biod, mountd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: exports

# nfsstat

## NFS statistics

Displays statistical information about the Network File System (NFS) and Remote Procedure Call (RPC). Nfsstat can also be used to reinitialize this information.

---

## /usr/etc/nfsstat [-cnrsz]

---

### ARGUMENTS

The options default to -csnr, that is, print everything and reinitialize nothing.

**-c** displays client information. Only the client side NFS and RPC information is printed. Can be combined with the -n and -r options to print client NFS or client RPC information only.

**-n** displays NFS information. NFS information for both the client and server side is printed. Can be combined with the -c and -s options to print client or server NFS information only.

**-r** displays RPC information. Works like the -n option above.

**-s** displays server information. Works like the -c option above.

**-z** zeros (reinitializes) statistics. This option is for use by the superuser only, and can be combined with any of the above options to zero particular sets of statistics after printing them.

### DESCRIPTION

The server RPC display includes the fields:

| | |
|---|---|
| *calls* | total number of RPC calls received |
| *badcalls* | total number of calls rejected |
| *nullrecv* | number of times no RPC packet was available when trying to receive |
| *badlen* | number of packets that were too short |
| *xdrcall* | number of packets that had a malformed header |

The server NFS display shows the number of NFS calls received (*calls*) and rejected (*badcalls*), and the counts and percentages for the various calls that were made.

The client RPC display includes the fields:

| | |
|---|---|
| *calls* | total number of RPC calls sent |
| *badcalls* | total of calls rejected by a server |
| *retrans* | number of times a call had to be retransmitted |
| *badxid* | number of times a reply did not match the call |
| *timeout* | number of times a call timed out |
| *wait* | number of times a call had to wait on a busy handle |
| *newcred* | number of times authentication information had to be refreshed |

The client NFS display shows the number of calls sent and rejected, as well as the number of times a handle was received (*nclget*), the number of times a call had to sleep while awaiting a handle (*nclsleep*), as well as a count of the various calls and their respective percentages.

## FILES

| | |
|---|---|
| /vmunix | system namelist |
| /dev/kmem | kernel memory |

## SEE ALSO

*Concentrix Network File System Manual*

# nslookup

## query Internet name servers interactively

The nslookup program queries Internet domain name servers. It has two modes: interactive and noninteractive. Interactive mode allows the user to query name servers for information about various hosts and domains or to print a list of hosts in a domain. Noninteractive mode is used to print just the name and requested information for a host or domain.

---

**nslookup** [*-option...* ] [ *host-to-find* | - [ *server* ]]

---

### DESCRIPTION

Interactive mode is entered in the following cases:

- when no arguments are given (the default name server will be used)
- when the first argument is a hyphen (-) and the second argument is the host name or Internet address of a name server

Noninteractive mode is used when the name or Internet address of the host to be looked up is given as the first argument. The optional second argument specifies the host name or address of a name server.

You can can specify the options listed under the set command below in the .nslookuprc file in the user's home directory if they are listed one per line. Options can also be specified on the command line if they precede the arguments and are prefixed with a hyphen. For example, to change the default query type to host information, and the initial timeout to 10 seconds, type:

```
nslookup -query=hinfo -timeout=10
```

Commands may be interrupted at any time by typing a CTRL C. To exit, type a CTRL D (EOF) or type exit. The command line length must be less than 256 characters. To treat a built-in command as a host name, precede it with an escape character (\). An unrecognized command will be interpreted as a host name.

host [*server*] Look up information for host using the current default server or using *server* if specified. If *host* is an Internet address and the query type is A or PTR, the name of the host is returned. If *host* is a name and does not have a trailing period, the default domain name is appended to the name. (This behavior depends on the state of the set options *domain, srchlist, defname,* and *search*). To look up a host not in the current domain, append a period to the name.

server or lserver *domain*
Change the default server to *domain*. lserver uses the initial server to look up information about *domain* while server uses the current default server. If an authoritative answer can't be found, the names of servers that might have the answer are returned.

root — Changes the default server to the server for the root of the domain name space. Currently, the host ns.nic.ddn.mil is used. (This command is a synonym for lserver ns.nic.ddn.mil.) The name of the root server can be changed with the set root command.

finger [*name*] [> *filename*]

finger [*name*] [>> *filename*]
Connects with the finger server on the current host. The current host is defined when a previous lookup for a host was successful and returned address information (see the set querytype=A command). *Name* is optional. > and >> can be used to redirect output in the usual manner.

ls *[option] domain* [> *filename*]

ls [*option] domain* [>> *filename*]
List the information available for *domain*, optionally creating or appending to *filename*. The default output contains host names and their Internet addresses. *Option* includes:

- *–t querytype* — lists all records of the specified type (see *querytype* below).
- *–a* — lists aliases of hosts in the domain. Synonym for *–t* CNAME.
- *–d* — lists all records for the domain. Synonym for *–t* ANY.
- *–h* — lists CPU and operating system information for the domain. Synonym for *–t* HINFO.
- *–s l* — lists well-known services of hosts in the domain. Synonym for *–t* WKS. When output is directed to a file, hash marks are printed for every 50 records received from the server.

view *filename* — Sorts and lists the output of previous ls command(s) with more.

help or ? — Prints a brief summary of commands.

exit — Exits the program.

set **keyword**[=*value*]
Use this command to change state information that affects the lookups. Valid keywords are:

- **all** — Prints the current values of the frequently-used options to set. Information about the current default server and host is also printed.
- **class=***value* Change the query class to one of:
  - IN — the Internet class.
  - CHAOS — the Chaos class.
  - HESIOD — the MIT Athena Hesiod class.
  - ANY — wildcard (any of the above). The class specifies the protocol group of the information. (Default = IN, abbreviation = cl)

**[no]debug** Turns debugging mode on. A lot more information is printed about the packet sent to the server and the resulting answer.
(Default = nodebug, abbreviation = [no]deb)

**[no]d2** Turns exhaustive debugging mode on. Essentially all fields of every packet are printed.
(Default = nod2)

**domain=***name*
Changes the default domain name to *name*. The default domain name is appended to a lookup request depending on the state of the *defname* and *search* options. The domain search list contains the parents of the default domain if it has at least two components in its name. For example, if the default domain is CC.Berkeley.EDU, the search list is CC.Berkeley.EDU and Berkeley.EDU. Use the set srchlist command to specify a different list. Use the set all command to display the list.
(Default = *value* from hostname, /etc/resolv.conf or LOCALDOMAIN, abbreviation = do)

**srchlist=***name1/name2/...*
Changes the default domain name to *name1* and the domain search list to *name1*, *name2*, etc. A maximum of 6 names separated by slashes (/) can be specified. For example:

```
set srchlist=lcs.MIT.EDU/ai.MIT.EDU/MIT.EDU
```

sets the domain to lcs.MIT.EDU and the search list to the three names. This command overrides the default domain name and search list of the set domain command. Use the set all command to display the list.

(Default = value based on *hostname*, /etc/resolv.conf or LOCALDOMAIN, abbreviation = srchl)

**[no]defname**
If set, appends the default domain name to a single-component lookup request (i.e., one that does not contain a period).
(Default = defname, abbreviation = [no]def)

**[no]search**
If the lookup request contains at least one period but doesn't end with a trailing period, appends the domain names in the domain search list to the request until an answer is received.
(Default = search, abbreviation = [no]sea)

**port=***value*
Changes the default TCP/UDP name server port to *value*.
(Default = 53, abbreviation = po)

**querytype=***value* or **type=***value*
Changes the type of information query to one of:

A the host's Internet address.

CNAME
the canonical name for an alias.

HINFO
the host CPU and operating system type.

MINFO
the mailbox or mail list information.

MX the mail exchanger.

NS the name server for the named zone.

PTR the host name if the query is an Internet address, otherwise the pointer to other information.

SOA the domain's "start-of-authority" information.

TXT the text information.

UINFO
the user information.

WKS the supported well-known services. Other types (ANY, AXFR, MB, MD, MF, NULL) are described in the RFC-1035 document. (Default = A, abbreviations = q, ty)

**[no]recurse**
Tells the name server to query other servers if it does not have the information. (Default = recurse, abbreviation = [no]rec)

**retry=*number***
Sets the number of retries to ***number***. When a reply to a request is not received within a certain amount of time (changed with set timeout), the timeout period is doubled and the request is resent. The retry value controls how many times a request is resent before giving up. (Default = 4, abbreviation = ret)

**root=*host***
Changes the name of the root server to ***host***. This affects the root command. (Default = ns.nic.ddn.mil., abbreviation = ro)

**timeout=*number***
Changes the initial timeout interval for waiting for a reply to ***number*** seconds. Each retry doubles the timeout period.
(Default = 5 seconds, abbreviation = ti)

**[no]vc** Always use a virtual circuit when sending requests to the server.
(Default = novc, abbreviation = [no]v)

**[no]ignoretc**
Ignore packet truncation errors.
(Default = noignoretc, abbreviation = [no]ig)

## DIAGNOSTICS

If the lookup request was not successful, an error message is printed. Possible errors are:

`Timed out`
The server did not respond to a request after a certain amount of time (changed with set timeout=*value*) and a certain number of retries (changed with set retry=*value*).

`No response from server`
No name server is running on the server machine.

`No records`
The server does not have resource records of the current query type for the host, although the host name is valid. The query type is specified with the set querytype command.

`Non-existent domain`
The host or domain name does not exist.

`Connection refused`

`Network is unreachable`
The connection to the name or finger server could not be made at the current time. This error commonly occurs with ls and finger requests.

`Server failure`
The name server found an internal inconsistency in its database and could not return a valid answer.

`Refused` The name server refused to service the request.

`Format error`
The name server found that the request packet was not in the proper format. It may indicate an error in nslookup.

## FILES

| | |
|---|---|
| /etc/resolv.conf | initial domain name and name server addresses. |
| $HOME/.nslookuprc | user's initial options. |
| /usr/share/misc/nslookup.help | summary of commands. |
| HOSTALIASES | file containing host aliases. |
| LOCALDOMAIN | overrides default domain. |

## SEE ALSO

named
*Concentrix-2800 System Reference: System Calls and C Library Routines:* resolver
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* resolver
RFC-1034 "Domain Names - Concepts and Facilities"
RFC-1035 "Domain Names - Implementation and Specification"

# ntalkd

## server for talk program

Listens at the UDP port used by the talk service as specified in /etc/services. Supports a conversation between two users on a TCP connection.

The ntalkd daemon is started as needed by /etc/inetd.

---

## /etc/ntalkd

---

### DESCRIPTION

Ntalkd is the server that notifies a user that somebody else wants to initiate a conversation. It acts a repository of invitations, responding to requests by clients wishing to rendezvous to hold a conversation.

In normal operation, a client, the caller, initiates a rendezvous by sending a CTL_MSG to the server of type LOOK_UP (see <protocols/talkd.h>). This action causes the server to search its invitation tables to see if an invitation currently exists for the caller (to speak to the callee specified in the message). If the lookup fails, the caller then sends an ANNOUNCE message causing the server to broadcast an announcement on the callee's login ports requesting contact. When the callee responds, the local server uses the recorded invitation to respond with the appropriate rendezvous address, and the caller and callee client programs establish a stream connection through which the conversation takes place.

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: talk, write
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services

# opt_software

## select software from a distribution tape

Reads the XDR-format table of contents found on software distribution tapes intended for use by diskless/NFS servers and allows the user to construct an extraction list of optional software.

---

**/usr/etc/install/opt_software** *arch tapedev* [*tapehost*]

---

After the table of contents is read, the user is presented with a series of prompts asking whether to select a particular software package from the distribution tape. In addition to the name, the user is told the size of the package in bytes along with an indication of the importance of the selection (required, desirable, common, or optional). The user selects the software by entering yes or no to each of the prompts as they are presented. The selection list is left in the file /tmp/EXTRACTLIST.<architecture> for use by setup_exec.

### ARGUMENTS

*arch* is the name of the software architecture, such as sun2, sun3, or sun4.

*tapedev* is the full pathname of the device where the tape is mounted. This program assumes that the table of contents is located in the second file on the tape. Alternatively, *tapedev* can be the pathname of an arbitrary file that contains the XDR format table of contents, which is assumed if the pathname does not start with /dev.

*tapehost* is specified if the tape device is located on a remote system. The presence of this option requires that *tapedev* be the pathname of an actual tape device.

### FILES

/tmp/EXTRACTLIST.<architecture>

### SEE ALSO

INSTALL, setup_exec

# pac

## printer/plotter accounting information

Reads the printer/plotter accounting files, accumulating the number of pages (the usual case) or feet (for raster devices) of paper consumed by each user, and printing out how much each user consumed in pages or feet and dollars.

---

**/etc/pac** [*option* ...] [*name* ...]

---

### ARGUMENTS

**-c** sorts the output by cost (default is alphabetically by name).

**-m** ignores the host name so that charges for a user with accounts on multiple machines are grouped together.

**-P***printer* enables accounting for the named printer. Normally, accounting is done for the default printer (site dependent) or the value of the environment variable `PRINTER` is used.

**-p***price* specifies the cost in dollars (default *price* is 0.02 or as specified in /etc/printcap).

**-r** reverses the sorting order.

**-s** summarized the accounting information in the summary accounting file; this is necessary since on a busy system, the accounting file can grow by several lines per day.

*name* print statistics only for the specified users (by default, statistics are printed for every user who has used any paper.)

### FILES

| | |
|---|---|
| /usr/adm/?acct | raw accounting files |
| /usr/adm/?_sum | summary accounting files |

### NOTES

- The relationship between the computed price and reality is as yet unknown.

# patchunix

## patch Concentrix-2800 image

Reads input files for lines that contain patch directives and/or adb commands and uses these directives/commands to patch a Concentrix-2800 image file and update its versionpatch value.

---

**/usr/alliant/patchunix [-l] [-u *unixfile*] [*file...*]**

---

### ARGUMENTS

**-l** requests a listing of patches in the Concentrix-2800 image.

**-u** *unixfile* specifies the file containing the Concentrix-2800 image; defaults to /unix.

*file* specifies a patch file to be scanned for input.

### DESCRIPTION

Patch files are processed in order. A patch file , which resembles a csh command script, consists of lines of the following types:

- **comment** – a blank line or a line beginning with #.
- **directive** – a line with the form *#keyword keyvalue*.
- **adb command** – signals the start of patch commands.
- **patch commands** – processed by adb.
- **end-of-input mark** – all lines after the end-of-input mark are ignored.

The directives are as follows:

**#pu_pnum** *integer*
range: 1–32 decimal, patch identifier (required).

**#pu_vmin** *hex_value*
range: 0x00000000 – 0xFFF00000, minimum version (inclusive) of Concentrix required to install patch (optional: default = 0x0).

**#pu_vmax** *hex_value*
range: pu_vmin – 0xFFFFFFFF, maximum version (exclusive) of Concentrix required to install patch (optional: default = 0xFFFFFFFF).

**#pu_chkl** *sym_name sym_hex_value*
compares the longword at location *sym_name* in the Concentrix-2800 image with *sym_hex_value*. If the comparison is not equal, the patch is not installed. You can include any number of pu_chkl directives (including none), but all must be satisfied.

## EXAMPLES

Sample patch file:

```
#!/bin/csh -f
#
#       Patch to fix "kmem_alloc" panic caused by
#       a failure to free kmem space if pathname characters
#       have the high bit set
#
#pu_pnum     1
#pu_vmin     0x04000000
#pu_vmax     0x04200000
#pu_chkl     pn_get+0x94  0x02400080
#
adb -w /unix /dev/kmem <<EOP
@patchtptr?w 0x0240
.+2?w 0x0080
.+2?w 0x6716
.+2?w 0x4855
.+2?w 0x204F
.+2?w 0x4878
.+2?w 0x0001
.+2?w 0x4EB9
.+2?W _pn_free
.+4?w 0x508F
.+2?w 0x4EF9
.+2?W _pn_get+0x9A
.+4?w 0x4EF9
.+2?W _pn_get+0x9E
pn_get+0x94?w       0x4EF9
.+2?W @patchtptr
patchtptr?W @patchtptr+0x22
EOP
#
```

List patches in the default Concentrix-2800 image:

```
# patchunix -l
```

Make a new /unix, applying patch in /usr/alliant/patches/fixbug-1.v5

```
# mv /unix /unix.prepatch
# cp /unix.prepatch /unix
# patchunix /usr/alliant/patches/fixbug-1.v5
```

## NOTES

- This utility is primarily for use by Customer Service personnel.
- See the administration procedure Patch_kernel for a list of kernel procedures and additional information on patching Concentrix-2800.
- The user must have read and write access to the Concentrix-2800 image to be patched, or read access if –l is specified.
- The characters 0x must be prepended to hexadecimal values in patch directives.
- The adb utility must be executable by the user.
- Diagnostix images are not patchable.

## SEE ALSO

Appendix B: Patch_kernel
*Concentrix-2800 Commands and Applications Manual*: adb
*Concentrix-2800 System Programming Guide*: Adb – Assembly Language Debugger

# ping

## send icmp echo packet (network debugging)

Repeatedly sends an icmp echo packet to a specified host system and reports whether or not a reply is received.

---

**/etc/ping [-r] [-v]** *host* [*packet-size*] [*count*]

---

### ARGUMENTS

| | |
|---|---|
| **-r** | bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it (for example, interface was dropped by routed). |
| **-v** | Verbose output. ICMP packets other than ECHO_RESPONSE that are received are listed. |
| *host* | host name or its Internet address. |
| *packet-size* | number of data bytes in the packet; defaults to 56. The total length of the packet is *packet-size* + 8. |
| *count* | number of packets. If you do not specify *count*, type CTRL C to stop the display. |

### DESCRIPTION

The DARPA Internet is a large and complex aggregation of network hardware, connected together by gateways. Tracking a single-point hardware or software failure is often difficult. Ping uses the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. ECHO_REQUEST datagrams ("pings") have an IP and ICMP header, followed by a struct timeval and an arbitrary number of pad bytes to fill out the packet.

For fault isolation, first run ping on the local host to verify that the local network interface is up. Then run ping on hosts and gateways further and further away. Ping sends one datagram per second, and prints one line of output for every ECHO_RESPONSE returned. No output is produced for no response. Round-trip times and packet loss statistics are computed. At program termination, a brief summary is displayed.

This program is intended for use in network testing, measurement, and management and should be used primarily for manual fault isolation. Because of the load that could be imposed on the network, use of ping during normal operations or from automated scripts is not recommended.

### SEE ALSO

ifconfig
*Concentrix-2800 Commands and Applications Manual*: netstat
*Concentrix-2800 System Reference: Miscellaneous* Routines, *Special Files, and File Formats:* icmp

# portmap

## DARPA port to RPC program number mapper

Converts RPC program numbers into DARPA protocol port numbers. This server must be running in order to make RPC calls.

The portmap daemon is typically invoked by /etc/rc.local.

---

## /usr/etc/rpc.portmap

---

### DESCRIPTION

When an RPC server is started, it tells portmap what port number it is listening to and what RPC program numbers it is prepared to serve. When a client wishes to make an RPC call to a given program number, it first contacts portmap on the server machine to determine the port number where RPC packets should be sent.

Portmap must be started before /etc/inetd or any other servers.

### NOTES

- If portmap crashes, all servers must be restarted.

### SEE ALSO

rpcinfo, inetd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: inetd.conf, servers

# ptynum

## change nptys flag in kernel image

Sets the value of the nptys flag; this flag determines the maximum number of pseudo terminals. The default is 64 pseudo terminals.

---

**/usr/alliant/ptynum** [*value* [*kernel-name*]]

---

### ARGUMENTS

*value* a decimal number that specifies the number of pseudo terminals. The maximum permissible value is 256. If you omit *value*, ptynum prints out the current value of nptys.

*kernel-name*
(defaults to /unix) name of the kernel image.

### EXAMPLES

The following example sets the maximum number of pseudo terminals to 80:

```
# /usr/alliant/ptynum 80
```

### NOTES

- Because ptynum patches a kernel variable, you must rerun it whenever you install a new kernel.

### SEE ALSO

*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: pty

# quot

## summarize file system ownership

Prints the number of blocks currently owned by each user in the named file system. If no filesystem is specified, a default name is assumed.

---

**/usr/etc/quot** [*option ...*] [*filesystem*]

---

### ARGUMENTS

**-a** generates a report for all mounted file systems.

**-c** prints three columns giving file size in blocks, number of files of that size, and cumulative total of blocks in that size file or smaller.

**-f** also prints the number of files owned by each user. Blocks are in column 1; number of files are in column 2.

**-h** estimates the number of blocks (does not account for files with holes).

**-n** runs the pipeline ncheck filesystem | sort +0n | quot -n filesystem to produce a list of all files and their owners.(not supported).

**-v** prints three columns giving the number of blocks not accessed in the last 90, 60, and 30 days.

*filesystem* device name of a file system; defaults to all mounted file systems.

### EXAMPLES

The following example prints file system ownership information for /dev/xdv11s:

```
# /usr/etc/quot /dev/xdv11s
```

### DIAGNOSTICS

- quot: read error at block 16 – Usually means an invalid file system name.

### FILES

| | |
|---|---|
| /etc/mtab | mounted file systems |
| /etc/passwd | user names |

### SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: ls, du

# quotacheck

## file system quota consistency checker

Examines each file system, builds a table of current disk usage, and compares this table against that stored in the disk quota file for the file system. If any inconsistencies are detected, both the quota file and the current system copy of the incorrect quotas are updated (the latter only occurs if an active file system is checked).

---

**/etc/quotacheck [-v] [-p] [*filesystem*]**

**/etc/quotacheck [-v] [-p] -a**

---

### ARGUMENTS

**-a** checks all file systems in /etc/fstab indicated as read-write with disk quotas.

**-p** runs parallel passes on the filesystems required, using the pass numbers in /etc/fstab in an identical fashion to fsck.

**-v** indicates the calculated disk quotas for each user on a particular file system.

*filesystem* device name of a file system; defaults to all mounted file systems.

### DESCRIPTION

Quotacheck is normally run at boot time from the /etc/rc file before enabling disk quotas with quotaon.

> **Note:** The distributed /etc/rc file has quotacheck commented out. You must uncomment it if you want it to run at boot time.

Quotacheck accesses the raw device in calculating the actual disk usage for each user. Thus, the file systems checked should be quiescent while quotacheck is running.

Quotacheck expects each file system checked to have a quota file named quotas in the root directory. If none is present, quotacheck ignores the file system.

Normally quotacheck reports only those quotas modified.

### FILES

/etc/fstab default file systems

### SEE ALSO

fsck, quotaon, rc
*Concentrix-2800 System Reference: System Calls and C Library Routines*: quotactl

# quotaon quotaoff

## turn file system quotas on and off

Enables and disables disk quotas. The file systems specified must be mounted at the time. The file system quota file must be present in the root directory of each specified file system and must be named quotas.

---

**/etc/quotaon [-v]** ***filesystem*** **...**

**/etc/quotaon [-v] -a**

**/etc/quotaoff [-v]** ***filesystem*** **...**

**/etc/quotaoff [-v] -a**

---

### ARGUMENTS

**-a** turns on/off quotas of all file systems in /etc/fstab indicated as read-write with disk quotas. This option is normally used at boot time to enable quotas.

**-v** displays a message for each file system as quotas are turned on/off.

***filesystem*** device name of a file system.

### DESCRIPTION

These commands update the status field of devices located in /etc/mtab to indicate when quotas are on or off for each file system.

The script /etc/rc.local contains the following lines:

```
# Handle quotas (Remove leading #'s from each of the four lines
#                following the next blank line when activating
#                quotas for your system.)

#                                  echo -n 'checking quotas: '     >/dev/console
#/etc/quotacheck -a -p                                            >/dev/console 2>
&1
#                                  echo  'done.'                   >/dev/console
#/etc/quotaon -a                                                  >/dev/console 2>
&1
```

To turn quotas on at boot time, uncomment the quotacheck and quotaon commands. For all file systems in which quotas are being enforced, ensure (1) that its root directory contains a file named quotas and (2) that its mount entry in /etc/fstab specfies quota as an option.

### EXAMPLE

```
# quotaon -v /usr
/usr: quotas turned on
```

## FILES

| | |
|---|---|
| quotas | quota file at the file system root |
| /etc/mtab | mounted file system table |
| /etc/fstab | static file system mounting table |

## SEE ALSO

edquota, quotacheck, repquota
*Concentrix-2800 Command and Applications*: quota
*Concentrix-2800 System Reference: System Calls and C Library Routines*: quotactl, rquota
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* mtab, fstab

# rc

## command scripts for starting system

Checks file systems and performs various system initialization tasks. These command scripts are invoked during the boot sequence.

---

**/etc/rc.diag**

**/etc/rc**

**/etc/rc.local**

---

### DESCRIPTION

Three scripts are invoked automatically during the boot sequence:

- **rc.diag** – Checks the Diagnostix-2800 root file system (/dev/xdv00p, by default) by invoking fsck with the -p option, mounts the file system, and invokes the script /console/console_loop. The script console.loop tests the FX/2800 hardware by invoking the diagnostic monitor; loads Concentrix-2800 under debugger control by invoking the FX/2800 booter; connects the console to Concentrix-2800 through tip (connect to a remote system); and responds to superuser requests if the superuser stops in FX/2800 startup mode or disconnects from Concentrix-2800. The script rc.diag is invoked from init when Diagnostix-2800 starts up.

- **rc** – Checks all files systems by invoking fsck with the -p option, unless a fastboot is in progress; makes a node for the swapping partition; recovers the password file if necessary; removes the nologin file; mounts all file systems; and invokes rc.local. Recovers files edited with vi; clears /tmp except for directories; starts the update, clock, and accounting daemons; starts networking daemons; starts the printer daemon. This script is invoked from init when Concentrix-2800 starts up.

- **rc.local** – Sets the hostname; starts the network, the network information service, and the network file system if appropriate; initializes the message-of-the-day file; starts the DARPA TELNET server, the DARPA internet file transfer protocol server, the talk daemon, the sendmail daemon, and the batch daemon, as appropriate; and performs other site-specific actions. This script is invoked from rc.

You must make the following modifications to /etc/rc.local as distributed:

- Replace myname with the name of your system on the hostname myname line.

- If your system is part of a domain, uncomment the domainname mydomain line and replace mydomain with the name of your domain.

- Modify the first ifconfig line to suit your network configuration. If you are not running subnets, you can delete netmask my-netmask and broacast my-broadcast. At any rate, ensure that my-broadcast identifies your network, not your host (the last decimal position must be 0). Delete this line if you have no Ethernet controllers. If you have more than one Ethernet controller, specify additional ifconfig lines. See ifconfig. Do not modify the second ifconfig line (for localhost).

- If your system uses the NIS, uncomment the 6 lines concerning ypserv and ypbind.

- Uncomment the lines pertaining to routed, named, timed, nfsd, biod, statd, and lockd, as appropriate.

At the end of /etc/rc.local, place any site-specific system startup procedures, for example, to set the scheduling vector. Use the full pathname of commands, for example, type /usr/alliant/setsched not setsched.

## EXAMPLES

Not intended for user invocation.

## FILES

| | |
|---|---|
| /etc/fstab | static information about file systems |
| /etc/mkstripe | striped partition database |
| /etc/motd | message-of-the-day file |
| /etc/mtab | mounted file system table |
| /etc/nologin | file whose existence prevents logins |
| /etc/passwd | password file |
| /fastboot | empty file whose existence signals a fastboot |
| /tmp | directory for temporary user files |

## SEE ALSO

accton, batchd, comsat, cron, fastboot, fsck, ftpd, inetd, init, lpd, mkstripe, mkswap, mount, portmap, reboot, rexecd, routed, rshd, rwhod, shutdown, swapon, telnetd, umount, update, ypbind, ypserv
*Concentrix-2800 Commands and Applications Manual*: hostname

# rdate

## set system data from a remote host

Rdate sets the local date and time from the hostname given as argument. You must be superuser on the local system. Typically rdate can be inserted as part of your /etc/rc.local startup script.

---

**/usr/ucb/rdate** *hostname*

---

### FILES

/etc/rc.local

### SEE ALSO

timed

# rdump

## file system dump across the network

Backs up a file system to magnetic tape or a disk file on a remote system.

---

**/etc/rdump** [*options* [*arguments*]] *filesystem*

---

### ARGUMENTS

Rdump is identical in operation to dump except that you must specify the **f** option and must specify its argument as *system:file*, where *system* is the name of the remote system and *file* is the name of the tape device or file on the remote system.

Since rdump is normally run by root, the name of the local machine must appear in the .rhosts file of the remote machine. If the filename argument is of the form user@machine:device, rdump attempts to execute as the specified user on the remote machine. The specified user must have an .rhosts file on the remote machine that allows root from the local machine. When rdump is called, the tape defaults to dumphost:/dev/rmt8. To direct the output to a desired remote machine, set up an alias for dumphost in the file /etc/hosts.

### DIAGNOSTICS

Same as dump with a few extra diagnostics related to the network.

### EXAMPLES

The following command performs a level 0 dump of the file system /dev/rxdv00f. If the dump completes successfully, /etc/dumpdates is updated. The output device is the Exabyte drive /dev/rschs (corresponds to f) on the system far.

```
# /etc/rdump 0ubdsf 32 6250 30000 far:/dev/rschs /dev/rxdv00f
```

### SEE ALSO

dump, rmt

# reboot

## reboot system

Reboots the system, or performs other actions depending on the arguments. Logs the reboot unless -l, -n, or -q is specified.

---

**/etc/reboot** [*argument ...*]

---

### ARGUMENTS

**-B** stops in VMEbus booter mode and prompts for names for the Diagnostix-2800 kernel, the Diagnostix-2800 kernel debugger, and the Diagnostix-2800 root file system. Type RETURN to continue with default file names. You can skip loading the Diagnostix-2800 debugger by typing ESCAPE. From the booter prompt, you can also type %h to get help and %ls to list the root directory. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-b** stops in FX/2800 booter mode and prompts for file names for the Concentrix-2800 kernel, the Concentrix-2800 kernel debugger, and the processor configuration file. Type RETURN to continue with default file names. From the booter prompt, you can also type ! (exclamation mark) to connect to Diagnostix-2800 (type CTRL d to return to the FX/2800 booter). This option causes only Concentrix-2800 to be rebooted.

**-c** same as -H.

**-D** stops in Alliant ddt mode and waits for commands. You can proceed by typing ESCAPE p or ESCAPE g. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-d** stops in Alliant ddt 2800 mode and waits for commands. You can proceed by typing ESCAPE p or ESCAPE g. This option causes only Concentrix-2800 to be rebooted.

**-H** halts the system (shuts the system down to console mode). Recommended for power-down. From the console, type u to exit from console mode and start the system. From a remote terminal, type o r RETURN u.

**-h** shuts the system down to Diagnostix-2800 mode. To start Concentrix-2800, type CTRL d, then respond y to the FX/2800 prompt.

**-l** inhibits the log message.

**-n** does not sync the disks (for emergencies only); inhibits the log message.

**-q** reboots quickly, without first shutting down running processes; inhibits the log message.

**-S** stops in Diagnostix-2800 mode. To continue the reboot, type CTRL d, then respond y to the FX/2800 prompt. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

-s stops at single-user mode in Concentrix-2800. To continue the reboot, type sync[RETURN], /etc/umount -a[RETURN], sync[RETURN], and [CTRL]d. This option causes only Concentrix-2800 to be rebooted.

## NOTES

- The command /etc/shutdown is preferred in an interactive environment in multi-user mode.
- The existence of the file /fastboot causes file system checking to be skipped.

## EXAMPLES

The following example reboots the system.

```
# /etc/reboot
```

The following example reboots the system and stops in Alliant ddt mode.

```
# /etc/reboot -D
```

The following example reboots the system and stops in Diagnostix-2800 mode.

```
# /etc/reboot -S
```

## FILES

| | |
|---|---|
| /dunix | Diagnostix-2800 kernel image |
| /fastboot | empty file whose existence signals a fastboot |
| /usr/adm/wtmp | login record of all logins and logouts |
| /unix | Concentrix-2800 kernel image |
| /vmunix | a symbolic link to /unix |

## SEE ALSO

fastboot, fsck, init, rc, shutdown

# renice

## alter priority of running processes

*Renice* alters the scheduling priority of one or more running processes.

---

**/etc/renice** *priority* [*option* [*argument*]] ...

---

### ARGUMENTS

*priority* specifies a numeric priority value.

**-p** *pid* specifies that subsequent arguments are process ID's (the default).

**-g** *pgrp* specifies that subsequent arguments are process group ID's. All processes in the process group have their scheduling priority altered.

**-u** *user* specifies that subsequent arguments are user names. All processes owned by the user have their scheduling priority altered.

### DESCRIPTION

Non-super-users can alter the priority only of processes they own, and can use *priority* values in the range 0 to PRIO_MAX (20). (This prevents abuse of system resources).

The superuser can alter the priority of any process and set the *priority* to any value in the range PRIO_MIN (-20) to PRIO_MAX. Useful priorities are: 20 (the affected processes run only when nothing else in the system wants to), 0 (the "base" scheduling priority), anything negative (to make things go very fast).

### EXAMPLE

The following command changes the priority of process ID's 987 and 32, and all processes owned by users daemon and root.

```
/etc/renice +1 987 -u daemon root -p 32
```

### FILES

/etc/passwd to map user names to user ID's

### NOTES

- Non super-users can not increase scheduling priorities of their own processes, even if they were the ones that decreased the priorities in the first place.

### SEE ALSO

*Concentrix-2800 System Reference: System Calls and C Library Routines*: getpriority, setpriority

# reportcfg

## report software/hardware configuration

Runs several configuration report utilities and examines Concentrix-2800 and layered products to generate a consolidated report of configuration dependencies on the system.

---

## /usr/alliant/reportcfg

---

The report includes the following information:

- Hardware configuration (syscfg -d)
- Hardware revision levels and serial numbers (soft console)
- Disk drive limits and label contents (kernel variables and diskf L)
- Microcode type and revision level (/console/ucode/ucode.type)
- Concentrix-2800 identification information (kernel variables)
- Paging and swapping space, memory restrictions, and other Concentrix-2800 configuration information (kernel variables)
- Layered products identification information (presence of various files)
- Contents of file system tables (/etc/fstab and /etc/stripetab)
- Disk usage (du)

### NOTES

- This utility is primarily for use by Customer Service personnel.
- The user needs read access to the Concentrix-2800 image (/unix and /dev/kmem).
- /console/bin/soft_console should have mode set-uid and be owned by root with staff as group-owner.
- The adb utility must be executable by the user.
- If standard locations are not used for layered products, they may not be found by this utility.

### SEE ALSO

diskf, syscfg
*Concentrix-2800 Commands and Applications*: adb

# repquota

## summarize quotas for a file system

Prints a summary of the disk usage and quotas for the specified file systems. For each user the current number of files and amount of space is printed, along with any quotas created with edquota. Only the superuser may view quotas with this utility.

---

**/etc/repquota [-v]** *filesystem ...*

**/etc/repquota [-v] -a**

---

### ARGUMENTS

**-a** reports on all file systems in /etc/fstab indicated as read-write with disk quotas.

**-v** reports all quotas, even if there is no usage.

*filesystem* device name of a file system; defaults to all mounted file systems.

### DESCRIPTION

Generally, disk usage and quotas are reported in kilobytes; these units are indicated by the suffix K. If this would cause the field to overflow, space is reported in megabytes with the suffix M, or in gigabytes with the suffix G.

### EXAMPLE

```
# repquota -v /usr
/dev/xdv00f (/usr):
                        Block limits                        File limits
User            used    soft    hard    timeleft     used    soft    hard    timeleft
bob      +-     11K     10K     20K    48.2 mins       6       0       0
judy     --      7K     10K     20K                    1       0       0
```

### FILES

| | |
|---|---|
| quotas | quota file at the file system root |
| /etc/fstab | default file systems |

### SEE ALSO

edquota, quotacheck, quotaon, quotaoff
*Concentrix-2800 Commands and Applications:* quota
*Concentrix-2800 System Reference: System Calls and C Library Routines*: quotactl, rquota
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* fstab

# restore

## incremental file system restore

Restores a file system or selected files from a file system from magnetic tape, a disk file, or standard input, written by the dump command. Restore has a shell-like interactive mode that allows you to examine the contents of a dump file and select specific files to restore.

> **NOTE:** Before invoking restore, set your current working directory to the root of the file system that was dumped. For example, if the tape contains the file system that was mounted on /usr, set your current working directory to /usr.

---

**/etc/restore** *options* [*filenames*]

---

### ARGUMENTS

The first argument is a list of options. *Options* must be specified as an unbroken character string containing one "Command" keyword and any number of "Modifier" keywords. Where modifiers take arguments (for example, **b** *factor*), the arguments follow the keyword string in corresponding order.

### Commands

**i** invokes interactive mode (see below).

**r** (full restore) restores the entire dump file into the current directory. Use this command carefully, for example, to restore a level 0 dump onto a clear file system or to restore an incremental dump immediately after restoring a level 0 dump. Restore leaves the file restoresymtable in the root directory of the file system to pass information between incremental restore passes. Remove the file when the last incremental tape is restored.

> **NOTE:** After a full restore, create a new level 0 dump tape. The old level 0 dump tape should not be used again on this file system. Although the contents of the files are unchanged, the restored directories have new i-nodes.

**R** requests a particular tape of a multi-volume set on which to restart a full restore (see the r command above). Allows restore to be interrupted and restarted.

**t** lists each specified *filename* if it occurs on the dump. If *filename* is not specified, the root directory (the entire dump) is listed, unless the h modifier is specified.

**x** extracts only specified files from the tape. The owner, modification time, and mode are restored (if possible) for each file.

*filenames*
specifies the files that are to be restored. A directory name refers to the files and (recursively) subdirectories of that directory unless the h modifier is specified (see below). Defaults to the entire contents of the dump file.

**Modifiers**

b *n* specifies the blocking factor for tape reads in even multiples of 1024 bytes. The blocking factor determines the number of bytes transferred per I/O request. A large blocking factor gives greater throughput, especially with cartridges, since larger transfers mean less time spent starting and stopping the tape drive and more time spent actually transferring data. The default blocking factor is 10 (10240-byte transfers).

NOTE: the blocking factor in the restore command must match the blocking factor used in the dump command. If you specify no blocking factor, restore determines the tapeblock size and uses it.

e causes restore to stop and issue a continue? query for all resync restore errors.

f specifies the device or file containing the dump tape. If the name of the file is –, restore reads from the standard input. Dump and restore can be used in a pipeline to copy a file system. The following example copies the /usr file system to /mnt:

```
# dump 0f - /usr | (cd /mnt; restore xf -)
```

h extracts a directory file, rather than the contents of the directory. Prevents recursive restoration of complete subtrees from the tape. This modifier is valid only with the i, t, and x commands.

m extracts by i-node number rather than by file name. Useful, if only a few files are being extracted, to avoid regenerating the complete pathname to the file. This modifier is valid only with the i and x commands.

s *n* skips to the *n*'th file when one tape holds multiple dump files. For example, the following command positions you at the fifth file on the tape:

```
# /etc/restore xbfs 32 /dev/rschs 5
```

u (update) overlays a new file system on top of an existing one by replacing hard or soft links with the contents of the dump tape. Normally, such replacement is prohibited by restore.

v (verbose) types the name of each file restored preceded by its file type. By default, restore works silently.

y does not ask whether to abort the restore if a tape error occurs. Restore tries to skip over the bad tape block(s) and continues as best it can.

## INTERACTIVE MODE

When you specify the i command, restore reads the directory information from the dump and provides a shell-like interface that allows you to move around the directory tree and select files to be extracted. The interactive commands are specified below. For commands that require an argument, the default is the current directory.

**ls** [*dir*] lists the current or specified directory. Entries that are directories are appended with a slash character. Entries that have been marked for extraction are prepended with an asterisk. If the verbose modifier is set, the i-node number of each entry is also listed.

**cd** *arg* changes the current working directory as specified.

**pwd** prints the full pathname of the current working directory.

**add** [*file*] adds the specified argument to the list of files to be extracted. If a directory is specified, it and all its descendents are added to the extraction list (unless the h modifier is specified on the command line). Files that are on the extraction list are prepended with an asterisk when they are listed by ls.

**delete** [*file*] deletes the specified argument from the list of files to be extracted. If a directory is specified, it and all its descendents are deleted from the extraction list (unless the h modifier is specified on the command line). The most expedient way to extract most of the files from a directory is to add the directory to the extraction list and then delete those files that are not needed.

**extract** extracts from the dump all the files that are on the list of files to be extracted. Restore asks which volume to mount. The fastest way to extract a few files is to start with the last volume and work towards the first volume.

**setmodes** sets the owners, modes, and times for all directories added to the extraction list; nothing is extracted. Useful for cleaning up after a restore is prematurely aborted.

**verbose** toggles the sense of the v modifer. When set, the verbose modifier causes the ls command to list the i-node numbers of all entries. Also prints information about each file as it is extracted.

**help** prints a summary of the available commands.

**quit** exits immediately, even if the extraction list contains files.

## DIAGNOSTICS

- Invalid key characters.
- Read errors. If y is in effect or you respond y, restore attempts to continue.
- If the dump extends over more than one tape volume, restore asks you to change tapes. If the x or i command is in effect, restore also asks which volume to mount. The fastest way to extract a few files is to start with the last volume and work towards the first volume.

Common errors resulting from consistency checks are as follows:

- *filename*: not found on tape – The specified file name is listed in the tape directory, but was not found on the tape. Caused by tape read errors while looking for the file, or from using a dump created on an active file system.
- expected next file *inumber*, got *inumber* – A file not listed in the directory was found. Caused by using a dump created on an active file system.
- Incremental tape too low – A tape that was written before the previous incremental tape or that has too low an incremental level has been loaded.
- Incremental tape too high – A tape that does not begin its coverage where the previous incremental tape left off or that has too high an incremental level has been loaded.
- Tape read error while restoring *filename* – A tape read error occurred. If a file name is specified, its contents are probably partially wrong. If an i-node is being skipped or the tape is trying to resynchronize, no extracted files are corrupted, although files may not be found on the dump tape.
- Tape read error while skipping over inode *i-node* – A tape read error occurred. If a file name is specified, its contents are probably partially wrong. If an i-node is being skipped or the tape is trying to resynchronize, no extracted files are corrupted, although files may not be found on the dump tape.
- Tape read error while trying to resynchronize – A tape read error occurred. If a file name is specified, its contents are probably partially wrong. If an i-node is being skipped or the tape is trying to resynchronize, no extracted files are corrupted, although files may not be found on the tape.

In addition, restore prints the error status returned by Concentrix-2800 for all tape errors.

## EXAMPLES

The following example restores the /dev/rxdv01h file system using the current working directory as the file system root:

```
# /etc/restore rbf 32 /dev/rschs /dev/rxdv01h
```

The following example prints the contents of the dump on /dev/rschs:

```
# /etc/restore tbf 32 /dev/rschs
```

The following example accesses the /dev/rxdv01h file system saved on /dev/rschs in interactive mode:

```
# /etc/restore ibf 32 /dev/rschs /dev/rxdv01h
```

## FILES

| | |
|---|---|
| /tmp/rstdir* | file containing directories on the tape |
| /tmp/rstmode* | owner, mode, and time stamps for directories |
| ./restoresymtable | information passed between incremental restores |

## NOTES

- Results are not guaranteed for dump tapes made on active file systems.

## SEE ALSO

dump, mount, makefs, rdump, rrestore

# rexd

## RPC-based remote execution server

Supports remote program execution.

The rexd server is started as needed by /etc/inetd.

---

## /usr/etc/rpc.rexd

---

### DESCRIPTION

This daemon is started by inetd whenever a remote execution request is made, if the following line is placed in /etc/servers:

```
rpc tcp /usr/etc/rpc.rexd 100017 1
```

For non-interactive programs, standard file descriptors are connected directly to TCP connections. Interactive programs involve pseudo-terminals, similar to the login sessions provided by rlogin. This daemon can use the NFS to mount file systems specified in the remote execution request.

### DIAGNOSTICS

Diagnostic messages are normally printed on the console, and returned to the requestor.

### FILES

| | |
|---|---|
| /dev/ttyp | pseudo-terminals used for interactive mode |
| /etc/passwd | authorized users |

### SEE ALSO

inetd
*Concentrix-2800 Commands and Applications Manual*: on
*Concentrix-2800 System Reference*: *System Calls and C Library Routines Routines*: rexd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: exports, servers, inetd.conf

# rexecd

## remote execution server

Provides remote execution facilities with authentication based on user names and encrypted passwords. Rexecd is the server for the rexec routine.

The rexd server is started as needed by /etc/inetd.

---

## /etc/rexecd

---

### DESCRIPTION

Rexecd listens for service requests at the port indicated in the exec service as specified in /etc/services. A service request initiates the following protocol:

1. The server reads characters from the socket up to a null (\0) byte. The resultant string is interpreted as an ASCII number, base 10.

2. If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the stderr. A second connection is then created to the specified port on the client's machine.

3. A null terminated user name of at most 16 characters is retrieved on the initial socket.

4. A null terminated, unencrypted password of at most 16 characters is retrieved on the initial socket.

5. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.

6. Rexecd then validates the user as with rlogind. If the authentication is successful, rexecd changes to the user's home directory, and establishes the user and group protections of the user. If any of these steps fail, the connection is aborted and a diagnostic message returned.

7. A null byte is returned on the initial socket and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by rexecd.

## DIAGNOSTICS

Except for the last one listed below, all diagnostic messages are returned on the initial socket, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 7 above upon successful completion of all the steps prior to the command execution).

- username too long – The name is longer than 16 characters.
- password too long – The password is longer than 16 characters.
- command too long – The command line passed exceeds the size of the argument list (as configured into the system).
- Login incorrect – No password file entry for the user name existed.
- Password incorrect – The wrong password was supplied.
- No remote directory – The *chdir* command to the home directory failed.
- Try again – A fork by the server failed.
- <shellname>: ... – The user's login shell could not be started. This error is retruned on the connection associated with stderr, and is not preceded by a flag byte.

## NOTES

- Indicating Login incorrect as opposed to Password incorrect is a security breach that allows people to probe a system for users with null passwords.

## SEE ALSO

*Concentrix-2800 System Reference*: *System Calls and C Library Routines Routines*: rexec
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services

# rlogind

## remote login server

Provides a remote login facility with authentication based on privileged port numbers from trusted hosts. Rlogind is the server for the rlogin program.

The rlogind daemon is started as needed by inetd.

---

## /etc/in.rlogind [-ln]

---

-l prevents any authentication based on the user's .rhosts file, unless the user is logging in as the superuser.

-n disables transport-level keepalive messages.

### DESCRIPTION

Rlogind listens for service requests at the port indicated by the login service as specified in the file /etc/services. When a service request is received, the following protocol is initiated:

1. The server checks the client's source port. If the port is not in the range 0–1023, the server aborts the connection.

2. The server checks the client's source address and requests the corresponding host name (see the C library routine gethostbyaddr, the file format for hosts, and the administration command named). If the host name cannot be determined, the dot notation representation of the host address is used.

Once the source port and address are checked, rlogind proceeds with the authentication process described in rshd. rlogind then allocates a pseudo terminal and manipulates file descriptors so that the slave half of the pseudo terminal becomes the stdin, stdout, and stderr for a login process. The login process is an instance of the login program invoked with the -f option if authentication has succeeded. If automatic authentication fails, the user is prompted to log in as on a standard terminal line. The -l option prevents any authentication based on the user's .rhosts file, unless the user is logging in as the superuser.

The parent of the login process manipulates the master side of the pseudo terminal, operating as an intermediary between the login process and the client instance of the rlogin program. In normal operation, the packet protocol described for pty is invoked to provide CTRL S/CTRL Q type facilities and propagate interrupt signals to the remote programs. The login process propagates the client terminal's baud rate and terminal type, as found in the environment variable TERM. The screen or window size of the terminal is requested from the client, and window size changes from the client are propagated to the pseudo terminal. Transport-level keepalive messages are enabled unless the -n option is present. The use of keepalive messages allows sessions to be timed out if the client crashes or becomes unreachable.

In the present version of the software, /etc/rlogind does not set the line discipline, control terminal, or process group of the current terminal before passing control to the login shell specified in /etc/passwd. Failure to make corresponding changes in user-provided shells results in the shell either hanging or aborting if it is started by rlogind (via an rlogin session).

If a shell hangs in this state, it has no control terminal (ps and mon report ? instead of a tty name).

While the Alliant-supplied shells /bin/csh and /bin/sh set up the terminal properly, they will fail to do so if exec'd by a user-supplied login shell program which does not, itself, set the line discipline. In this event, /bin/csh or /bin/sh hangs (with no control terminal).

The line discipline of a terminal must be "new" for the control terminal or process group to be set. Some login shell programs also need to establish the control terminal and process group. The following example does both:

```
#include <stdio.h>
#include <sgtty.h>
#include <sys/ioctl.h>

 char   *ttynam;              /* control terminal's name */
 int    input;                /* control terminal's file descriptor */
 int    shpgrp;               /* current process group / pid */
 int    ldisc;                /* line discipline */

   if (argc == 1 && *argv[0] == '-')
   {
        ldisc = NTTYdisk;
        input = fileno(stdin);
        (void) ioctl(input,TIOCSETD,&ldisc)

        if (isatty(input))
        {
            ttynam = (char *) ttyname(input);
            close(input);
            input = open(ttynam, O_RDWR);
            shpgrp = getpid();
            (void) ioctl(input, TIOCSPGRP, (char *)&shpgrp);
            (void) setpgrp(0,shpgrp);
        }
    }
```

Note that leaving the line discipline as "new," as in the example, may not be desired (particularly if it executes the Bourne shell).

## DIAGNOSTICS

All diagnostic messages are returned on the connection associated with the stderr, after which any network connections are closed. An error is indicated by a leading byte with a value of 1.

- Try again – A fork by the server failed.
- /bin/sh: ... – The user's login shell could not be started.

## NOTES

- The authentication procedure assumes the integrity of each client machine and the connecting medium. The procedure is insecure but useful in an "open" environment.
- A facility to allow all data exchanges to be encrypted should be present.
- A more extensible protocol should be used.

## SEE ALSO

rshd
*Concentrix-2800 System Reference*: *System Calls and C Routines*: ruserok
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: pty, services, hosts

# rmail

## handle remote mail received via uucp

rmail interprets incoming mail received through uucp, collapsing "From" lines in the form generated by binmail into a single line of the form return-path!sender, and passing the processed mail on to sendmail. rmail is explicitly designed for use with uucp and sendmail.

---

**rmail** *recipient*

---

### SEE ALSO

sendmail
*Concentrix-2800 Commands and Applications Manual*: binmail, uucp

# rmt

## remote magtape protocol module

Used by the remote dump and restore programs in manipulating a magnetic tape drive through an interprocess communication connection. Typically started by an rexec or rcmd call.

---

## /etc/rmt

---

### DESCRIPTION

The rmt program accepts requests specific to the manipulation of magnetic tapes, performs the commands, then responds with a status indication. All responses are in ASCII and in one of two forms. Successful commands have responses as follows, where *number* is an ASCII representation of a decimal number:

A*number*\n

Unsuccessful commands have responses as follows, where *error-number* is one of the possible error numbers described in the *Concentrix-2800 System Reference: System Calls and C Library Routines* and *error-message* is the corresponding error string as printed from a call to perror:

E*error-number*\n*error-message*\n

The protocol is comprised of the following commands (a space separates each token):

O *device mode*
: opens the specified *device* using the indicated *mode*. *Device* is a full pathname and *mode* is an ASCII representation of a decimal number suitable for passing to the open system call. If a device had already been opened, it is closed before a new open is performed.

C *device*
: closes the currently open device. The *device* specified is ignored.

L *whence offset*
: performs an lseek system call using the specified parameters. The response value is that returned from the lseek call.

W *count*
: writes data onto the open device. Rmt reads *count* bytes from the connection, aborting if a premature end-of-file is encountered. The response value is that returned from the write system call.

R *count*
: reads *count* bytes of data from the open device. If *count* exceeds the size of the data buffer (10 kilobytes), it is truncated to the data buffer size. Rmt then performs the requested read system call and responds with A*count-read*\n if the read is successful; otherwise an error in the standard format is returned. If the read is successful, the data read is then sent.

I *operation count*
: performs an MTIOCOP ioctl system call using the specified parameters. The parameters are interpreted as the ASCII representations of the decimal values to place in the mt_op and mt_count fields of the structure used in the ioctl call. The return value is the *count* parameter when the operation is successful.

S
: returns the status of the open device, as obtained with a MTIOCGET ioctl call. If the operation is successful, an acknowledgement is sent with the size of the status buffer, then the status buffer is sent (in binary).

Any other command causes rmt to exit.

## DIAGNOSTICS

All responses are of the form described above.

## NOTES

- People tempted to use this for a remote file access protocol are discouraged.

## SEE ALSO

rdump, rrestore
*Concentrix System Refereence Manual*: *System Calls and C Library Routines Routines*: rcmd, rexec, perror
*Concentrix System Refereence Manual: Miscellaneous Routines, Special Files, and File Formats*: mtio

# route

## manually manipulate the routing tables

Manually manipulates the network routing tables. Typically not needed, as the network routing daemon routed should tend to this task.

---

**/etc/route [-f] [-n] [*command args*]**

---

### ARGUMENTS

**-f** flushes the routing tables of all gateway entries. If used in conjunction with one of the commands, the tables are flushed prior to the command's application.

**-n** prevents attempts to print host and network names symbolically when reporting actions.

*command* is one of the following:

**add** adds a route.

**delete** deletes a route.

### DESCRIPTION

All commands have the following syntax:

*command* [**net**|**host**] *destination gateway* [*metric*]

*destination* destination host or network.

*gateway* next-hop gateway to which packets should be addressed.

*metric* is a count indicating the number of hops to the destination. The metric is required for add commands. The metric must be 0 for a directly attached network and non-zero if the route uses one or more gateways. If the metric is 0, the specified gateway is the address of this host on the common network, indicating the interface to be used for transmission.

The optional keywords net and host force the destination to be interpreted as a network or a host, respectively. Otherwise, the routes to a particular host are distinguished from those to a network by interpreting the Internet address associated with *destination*. If *destination* has a local address part of INADDR_ANY, or if *destination* is the symbolic name of a network, the route is assumed to be to a network; otherwise, it is assumed to be a route to a host. If the route is to a destination connected via a gateway, *metric* should be greater than 0.

All symbolic names specified for a *destination* or *gateway* are looked up first as a host name by using the C routine gethostbyname. If this lookup fails, the C routine getnetbyname is then used to interpret the name as that of a network.

Route uses a raw socket and the SIOCADDRT and SIOCDELRT ioctl's to do its work. Only the super-user can modify the routing tables.

## DIAGNOSTICS

- add [host|network] %s: gateway %s flags %x – The specified route is being added to the tables. The values printed are from the routing table entry supplied in the ioctl call. If the gateway address used is not the primary address of the gateway (the first one returned by gethostbyname), the gateway address is printed numerically as well as symbolically.
- delete [host|network] %s: gateway %s flags %x – As above, but when deleting an entry.
- %s %s done – When the -f flag is specified, each routing table entry deleted is indicated with a message of this form.
- Network is unreachable – An attempt to add a route failed because the gateway listed is not on a directly connected network. The next-hop gateway must be specified.
- not in table – A delete operation was attempted for an entry not present in the tables.
- routing table overflow – An add operation was attempted, but the system was low on resources and was unable to allocate memory to create the new entry.

## NOTES

- The change operation is not implemented. Add the new route, then delete the old one.

## SEE ALSO

routed
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: hosts, networks,

# routed

## network routing daemon

Manages the network routing tables. The routing daemon uses a variant of the Xerox NS Routing Information Protocol in maintaining up-to-date kernel routing table entries. The generalized protocol is capable of use with multiple address types, but is currently used only for Internet routing within a cluster of networks.

The routed daemon is typically invoked by /etc/rc.local. (As distributed, the routed invocation is commented out in /etc/rc.local. Uncomment the invocation if you want routed to start automatically with the system.)

---

**/etc/routed [-d] [-g] [-s] [-q] [-t] [*logfile*]**

---

### ARGUMENTS

**-d** enables additional debugging information to be logged, such as bad packets received.

**-g** on internetwork routers, offers a route to the default destination. Typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers.

**-q** is the opposite of the **-s** option.

**-s** forces routed to supply routing information whether it is acting as an internetwork router or not. This is the default if multiple network interfaces are present, or if a point-to-point link is in use.

**-t** prints all packets sent or received on the standard output. Routed does not divorce itself from the controlling terminal, so that interrupts from the keyboard kill the process.

*logfile* specifies the name of the file in which to log routed's actions. This log contains information about any changes to the routing tables and, if not tracing all packets, a history of recent messages sent and received that are related to the changed route.

### DESCRIPTION

In normal operation, routed listens on the udp socket for the route service for routing information packets. If the host is an internet router, routed periodically supplies copies of its routing tables to any directly connected hosts and networks.

When routed is started, it uses the SIOCGIFCONF ioctl to find those directly connected interfaces that are configured into the system and are marked up (the software loopback interface is ignored). If multiple interfaces are present, it is assumed that the host forwards packets between networks. Routed then transmits a request packet on each interface (using a broadcast packet if the interface supports it) and enters a loop, listening for request and response packets from other hosts.

When a request packet is received, routed formulates a reply based on the information maintained in its internal tables. The response packet generated contains a list of known routes, each marked with a "hop count" metric (a count of 16, or greater, is considered "infinite"). The metric associated with each route returned provides a metric *relative to the sender.*

Response packets received by routed are used to update the routing tables if one of the following conditions is satisfied:

1. No routing table entry exists for the destination network or host, and the metric indicates the destination is reachable (that is, the hop count is not infinite).

2. The source host of the packet is the same as the router in the existing routing table entry. That is, updated information is being received from the very internet router through which packets for the destination are being routed.

3. The existing entry in the routing table has not been updated for some time (defined to be 90 seconds) and the route is at least as cost effective as the current route.

4. The new route describes a shorter route to the destination than the one currently stored in the routing tables, determined by comparing the metric of the new route against the one stored in the table.

When an update is applied, routed records the change in its internal tables and updates the kernel routing table. The change is reflected in the next response packet sent.

In addition to processing incoming packets, routed periodically checks the routing table entries. If an entry has not been updated for three minutes, the entry's metric is set to infinity and marked for deletion. Deletions are delayed an additional 60 seconds to insure the invalidation is propagated throughout the local internet.

Hosts acting as internet routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks. The response is sent to the broadcast address on nets capable of that function, to the destination address on point-to-point links, and to the router's own address on other networks. The normal routing tables are bypassed when sending gratuitous responses. The reception of responses on each network is used to determine that the network and interface are functioning correctly. If no response is received on an interface, another route may be chosen to route around the interface, or the route may be dropped if no alternative is available.

In addition to the facilities described above, routed supports the notion of "distant" *passive* and *active* gateways. When routed is started, it reads the file /etc/gateways to find gateways that cannot be located using only information from the SIOGIFCONF ioctl. Gateways specified in this manner should be marked passive if they are not expected to exchange routing information, while gateways marked active should be willing to exchange routing information (that is, they should have a routed process running on the machine). Passive gateways are maintained in the routing tables forever and information regarding their existence is included in any routing information transmitted. Active gateways are treated equally to network interfaces. Routing information is distributed to the gateway and if no routing information is received for a period of the time, the associated route is deleted.

External gateways are also passive, but are not placed in the kernel routing table nor are they included inrouting updates. The function of external entries is to inform routed that another routing process will install such a route, and that alternate routes to that destination should not be installed. Such entries are only required when both routers may learn of routes to the same destination.

The file /etc/gateways is comprised of a series of lines, each in the following format:

< net|host > *name1* gateway *name2* metric *value* < passive|active|external >

The net or host keyword indicates if the route is to a network or a specific host.

*Name1* is the name of the destination network or host. This can be a symbolic name located in /etc/networks or /etc/hosts (or, if started after named, known to the name server), or an Internet address specified in "dot" notation; see inet.

*Name2* is the name or address of the gateway to which messages should be forwarded.

*Value* is a metric indicating the hop count to the destination host or network.

The keyword passive, active, external indicates if the gateway should be treated as passive or active (as described above), or whether the gateway is external to the scope of the routed protocol.

Internetwork routers that are directly attached to the Arpanet or Milnet should use the Exterior Gateway Protocol (EGP) to gather routing information rahter than using a static routing table of passive gateways. EGP is required in order to provide routes for local networks to the rest of the Internet system. Sites needing assistance with such configurations should contact the Computer Systems Research Group at Berkeley.

## FILES

| | |
|---|---|
| /etc/gateways | routing daemon data base |

## NOTES

- The kernel's routing tables may not correspond to those of routed for short periods of time while processes using existing routes exit.

## SEE ALSO

*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: inet, udp

# rpcinfo

## report RPC information

Makes an RPC call to an RPC server and reports the results.

---

**/usr/etc/rpcinfo [-p] [*host*]**

**/usr/etc/rpcinfo [-n *portnum*] [-u] *host program* [*version*]**

**/usr/etc/rpcinfo [-n *portnum*] [-t] *host program* [*version*]**

**/usr/etc/rpcinfo -b *program* [*version*]**

**/usr/etc/rpcinfo -d *program* [*version*]**

---

### ARGUMENTS

**-p** probes the portmapper on *host* and prints a list of all registered RPC programs. *Host* defaults to the value returned by hostname.

**-u** makes an RPC call to procedure 0 of *program* on the specified *host* using UDP, and reports whether a response was received.

**-t** makes an RPC call to procedure 0 of *program* on the *host* using TCP, and reports whether a response was received.

**-n** Use *portnum* as the port number for the **-t** and **-u** options instead of the port number given by the portmapper.

**-b** makes an RPC broadcast to procedure 0 of *program* and *version* using UDP and reports all hosts that respond.

**-d** deletes registration for the RPC service of the specified program and version. Only the superuser can exercise this option.

*program* the name or number of an RPC program.

*version* version of the program. If *version* is not specified, rpcinfo attempts to find all the registered version numbers for *program* by calling version 0 (which is presumed not to exist; it it does exist, rpcinfo calls an extremely high version number instead) and attempts to call each registered version.

**Note:** The version number is required for the **-b** and **-d** options.

### EXAMPLES

To show all of the RPC services registered on the local machine use::

```
example% rpcinfo -p
```

To show all of the RPC services registered on the machine klaxon use:

```
example% rpcinfo -p klaxon
```

To show all machines on the local net that are running the NIS use:

```
example% rpcinfo -b ypserv 'version' | uniq
```

where *version* is the current NIS version obtained from the results of the **-p** option shown above.

To delete the registration for version 1 of the walld service use:

```
example% rpcinfo -d walld 1
```

## FILES

| | |
|---|---|
| /etc/rpc | names for RPC program numbers |

## SEE ALSO

portmap
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats:* rpc
*Concentrix Network File System Manual:* chapters 2, 3, and 5.

# rquotad

## remote quota server

Returns quotas for a user of a local file system that is mounted by a remote machine over the NFS. The results are used by quota to display user quotas for remote file systems. The rquotad daemon is an rpc server that is normally invoked by inetd.

---

## /usr/etc/rpc.rquotad

---

### FILES

| | |
|---|---|
| quotas | quota file at the file system root |

### SEE ALSO

edquota, inetd, quotacheck, quotaon, quotaoff, repquota
*Concentrix-2800 Commands and Applications Manual*: quota
*Concentrix-2800 System Reference: System Calls and C Library Routines*: quotactl, rquota
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: fstab, nfs, rpc, services

# rrestore

## restore a file system dump across the network

Restores a file system or selected files from a file system from magnetic tape or a disk file on a remote system, written by the rdump (or dump) command. Rrestore has a shell-like interactive mode that allows you to examine the contents of a dump file and select specific files to restore.

> NOTE: Before invoking rrestore, set your current working directory to the root of the file system that was dumped. For example, if the tape contains the file system that was mounted on /usr, set your current working directory to /usr.

---

**/etc/rrestore** *options* [*filenames*]

---

### ARGUMENTS

Rrestore is identical in operation to restore except that you must specify the f option and must specify its argument as *system:file*, where *system* is the name of the remote system and *file* is the name of the tape device or file on the remote system.

Rrestore creates a remote server, /etc/rmt, on the client machine to access the tape device. If the file is specified as user@machine:device, rrestore attempts to execute as the specified user on the remote machine. The specified user must have an .rhosts file on the remote machine that allows root from the local machine.

### DIAGNOSTICS

Same as restore with a few extra diagnostics related to the network.

### EXAMPLES

The following example accesses the /dev/rxdv01h file system saved on the file /dev/rschs on the system far in interactive mode:

```
# /etc/restore ibf 32 far:/dev/rschs /dev/rxdv01h
```

### SEE ALSO

restore, rmt

# rshd

## remote shell server

Provides remote execution facilities with authentication based on privileged port numbers from trusted hosts. Rshd is the server for the rcmd routine and, consequently, for the rsh command.

The rshd daemon is started as needed by /etc/inetd.

---

## /etc/rshd

---

### DESCRIPTION

Rshd listens for service requests at the port indicated in the cmd service as specified in the file /etc/services. A service request initiates the following protocol:

1. The server checks the client's source port. If the port is not in the range 0–1023, the server aborts the connection.

2. The server reads characters from the socket up to a null (\0) byte. The resultant string is interpreted as an ASCII number, base 10.

3. If the number received in step 2 is non-zero, it is interpreted as the port number of a secondary stream to be used for the stderr. A second connection is then created to the specified port on the client's machine. The source port of this second connection is also in the range 0–1023.

4. The server checks the client's source address and requests the corresponding host name (see the C routine gethostbyaddr, the file format for hosts, and the administration command named). If the host name cannot be determined, the dot notation representation of the host address is used.

5. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity on the client's machine.

6. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity to use on the server's machine.

7. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.

8. Rshd validates the user according to the following steps. The local (server end) user name is looked up in the password file and a chdir is performed to the user's home directory. If either the lookup or chdir fail, the connection is terminated. If the user is not the super-user (user id 0), the file /etc/hosts.equiv is consulted for a list of trusted hosts. If the client's host name is present in this file, the authentication is successful. If the lookup fails, or the user is the super-user, the file .rhosts in the home directory of the remote user is checked for the machine name, the identity of the user on the client's machine, and correct access permission. If this lookup fails, or if the .rhosts file is writable by any other than the owner, the connection is terminated.

9. A null byte is returned on the initial socket and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by rshd.

## DIAGNOSTICS

Except for the last one listed below, all diagnostic messages are returned on the connection associated with the stderr, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 9 above upon successful completion of all the steps prior to the execution of the login shell).

- locuser too long – The name of the user on the client's machine is longer than 16 characters.
- remuser too long – The name of the user on the remote machine is longer than 16 characters.
- command too long – The command line passed exceeds the size of the argument list (as configured into the system).
- Login incorrect – No password file entry for the user name existed.
- No remote directory – The chdir command to the home directory failed.
- Permission denied – The authentication procedure described above failed.
- Can't make pipe – The pipe needed for the stderr, wasn't created.
- Try again – A fork by the server failed.
- <shellname>: ... – The user's login shell could not be started. This message is returned on the connection associated with stderr, and is not preceded by a flag byte.

## NOTES

- The authentication procedure used here assumes the integrity of each client machine and the connecting medium. The procedure is insecure but useful in an open environment.

## SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: rsh
*Concentrix-2800 System Reference*: System Calls and C Library Routines Routines: rcmd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services, hosts

# rstatd

## kernel statistics server

The rstatd command is a server that returns performance statistics obtained from the kernel. These statistics are graphically displayed by perfmeter (for example, on Sun systems).

The rstatd daemon is started as needed by /etc/inetd.

---

## /usr/etc/rpc.rstatd

---

### NOTES

Systems with disk drivers to be monitored by this daemon must be configured so as to report disk (_dk_xfer) statistics.

### SEE ALSO

inetd
*Concentrix-2800 Commands and Applications Manual:* perfmeter
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: services

# rtnamed

## logical name daemon

Stores name-to-object mappings in a set of tables (not to be confused with file system tables). Rtnamed is an rpc-based server that accepts requests from the privileged commands listed under "See Also" and from the Concentrix-2800 kernel.

The rtnamed daemon is typically started by the system start-up script /etc/rc.local.

---

## /usr/alliant/rtnamed

---

### DESCRIPTION

A name table holds mappings of object names to object identifiers. A table is identified by a null-terminated ascii string, and must be unique. Associated with each table are the uid and gid of the table creator, used in permission checking during every table access. Since only the super-user can create tables, the uid is always 0. The maximum and current number of table entries are maintained to control storage allocation. A protection attribute is associated with each table to control all accesses to the table, including name entry, lookup, and removal.

Whether the table is privileged or non-privileged is noted in the data base. A privileged table requires the client process to be a privileged process; a non-privileged process can only perform lookups. In either case, the protection mechanism is used, although the authentication parameters can easily be forged by the client.

An object name is a logical way of referring to an object in the system, and is identified by a null-terminated ascii string. The name must be unique within the table, that is, an object name can be the same as another object in a different table, but never the same table. Associated with an object name is an object identifier. This variable length, array of integers typically references a physical system object. Note that the naming daemon places no meaning or size on the object identifier, but only provides a protected mechanism for translating names to objects. Objects have a uid and gid that identify the user that entered the object into the table. Objects have protection that is specified when the object is entered. These attributes provide a finer grain of access control than the table protection. An object is purgable, that is, can be removed from all tables.

The logical name daemon provides two fundatmental capabilities:

- **Name translation** – Allows applications to be developed using logical name for physical components in the system without being coupled to the physical location. For example, mapping process virtual space to external bus space (via the rt_map_vtoe system call) can use logical name to determine the external bus space and decouple it from configuration changes or installation dependencies that may occur over the life of the application.

- **Object protection** – Allows applications to control access to the system objects, whether they be actual physical objects or system controlled data structures. For example, the use of real-time semaphores or events is protected by specifying table and object protections that limit lookups to a particular user or group.

## NOTES

- To limit access to the name daemon, a reserved port bind is attempted. This system call fails if the user is not the superuser. However, the registration function is not privileged and can easily be corrupted.

## FILES

<sys/rtname.h>
<multi/name_xdr.o>

## SEE ALSO

rtncrtbl, rtndeltbl, rtndumptb, rtnenter, rtnlookup, rtnremove, rtnpurge
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtncrtbl

## create logical name table

Creates a logical name table that maps textual names to arbitrary data. Only the super-user can create logical name tables (your real identification must be root).

---

**/usr/alliant/rtncrtbl** *tblname maxents priv* **[-w:[e][r][l]] [-g:[e][r][l]]**

---

### ARGUMENTS

*tblname* specifies the name of the table, which must be unique among logical name tables.

*maxents* specifies the maximum number of entries permitted in the table.

*priv* one of the following:

**0** non-privileged.

**1** privileged. A privileged table requires the client process to be a privileged process; a non-privileged process can only perform lookups.

**-w:[e][r][l]**
specifies access permissions for world, that is, processes whose effective uid and gid do not match that of the creator. Defaults to no permissions.

**e** permission to enter names into the table.

**l** permission to look up names in the table.

**r** permission to remove names from the table.

**-g:[e][r][l]**
specifies access permissions for group, that is, processes whose effective uid does not match that of the table creator, but whose effective gid does match. Defaults to no permissions.

**e** permission to enter names into the table.

**l** permission to look up names in the table.

**r** permission to remove names from the table.

## EXAMPLES

The following example creates the logical name table xtbl with access restricted to the owner (the super-user). The table can hold 16 entries.

```
# /usr/alliant/rtncrtbl xtbl 16 1
```

The following example creates the logical name table ytbl with as much access as possible. The table can hold 32 entries.

```
# /usr/alliant/rtncrtbl ytbl 32 0 -w:erl -g:erl
```

## NOTES

- Owner permissions are always complete access: enter, lookup, and remove, plus dump and delete.
- Group and world permissions do not include dump and delete.

## DIAGNOSTICS

- privilege violation – You must be super-user. (Rtncrtbl accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtncrtbl can only be run by super-user.)

## SEE ALSO

rtnamed, rtndeltbl, rtndumptbl, rtnenter, rtnlookup, rtnremove, rtnpurge
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtndeltbl

## delete logical name table

Deletes a specified logical name table or all tables. Only the uper-user can delete a logical name table.

---

**usr/alliant/rtndeltbl** [*tblname*]

---

### ARGUMENTS

*tblname* specifies the name of an existing logical name table. Defaults to all logical name tables.

### NOTES

This administrative command (like the system call name_deltbl) will delete an entire logical name table along with any contents. This can cause the references to permanent objects to be lost. If no other references to a permanent object remain, the object becomes inaccessible, but still uses space in internal system tables.

### EXAMPLES

The following example deletes the logical name table xtbl:

```
# /usr/alliant/rtndeltbl xtbl
```

### DIAGNOSTICS

- privilege violation – You must be super-user. (Rtndeltbl accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtndeltbl can only be run by super-user.)

### SEE ALSO

rtnamed, rtncrtbl, rtndumptbl, rtnenter, rtnlookup, rtnpurge, rtnremove
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtndumptbl

## dump logical name table

Displays the contents of a specified logical name table or all tables on standard output.

---

**/usr/alliant/rtndumptbl** [*tblname*]

---

### ARGUMENTS

*tblname* specifies the name of an existing logical name table. Defaults to all logical name tables.

### EXAMPLES

The following example displays the contents of the logical name table xtbl:

```
# /usr/alliant/rtndumptbl xtbl

Dump of name table: 'xtbl'
        uid: 0, gid: 10
        tblprot:
                owner: ENTER REMOVE LOOKUP DUMP DELETE
                group: ENTER REMOVE LOOKUP
                world:
        maxents: 9, curents: 0
        (privileged)
Hash chain 171
        name: x1, id: 0x1
        (purgeable)
        objprot:
                owner: ENTER REMOVE LOOKUP DUMP DELETE
                group:
                world:
Hash chain 173
        name: x3, id: 0x1 0x2 0x3
        objprot:
                owner: ENTER REMOVE LOOKUP DUMP DELETE
                group: REMOVE LOOKUP
                world: LOOKUP
```

### DIAGNOSTICS

- privilege violation – You must be super-user. (Rtndumptbl accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtndumptbl can only be run by super-user.)

### SEE ALSO

rtnamed, rtncrtbl, rtndeltbl, rtnenter, rtnlookup, rtnpurge, rtnremove
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtnenter

## enter logical name into table

Enters a name-to-object mapping into a logical name table.

---

**/usr/alliant/rtnenter** *tblname objname objid purflg* **[-w:[r][l]] [-g:[r][l]]**

---

### ARGUMENTS

*tblname* specifies the name of an existing logical name table.

*objname* specifies a name for the object. The combination of *tblname* and *objname* must be unique.

*objid* specifies the object as a list of hexadecimal numbers separated by white space and enclosed in single or double quotes. The quotes are not necessary if only one number is specified.

*purflg* one of the following values:

**0** does not allow purging of this name entry.

**1** allows purging of this name entry.

**-w:[r][l]**

specifies access permissions for world, that is, processes whose effective uid and gid do not match that of the creator. Defaults to no permissions.

**l** permission to look up the name.

**r** permission to remove the name.

**-g:[r][l]**

specifies access permissions for group, that is, processes whose effective uid does not match that of the table creator, but whose effective gid does match. Defaults to no permissions.

**l** permission to look up names in the table.

**r** permission to remove names from the table.

## EXAMPLES

The following example enters the name x1 into the logical name table xtbl:

```
# /usr/alliant/rtnenter xtbl x1 1 1
```

The following example enters the name x3 into the logical name table xtbl:

```
# /usr/alliant/rtnenter xtbl x3 "1 2 3" 0 -w:1 -g:rl
```

## DIAGNOSTICS

- privilege violation – You must be super-user. (Rtnenter accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtnenter can only be run by super-user.)

## SEE ALSO

rtnamed, rtncrtbl, rtndeltbl, rtndumptbl, rtnlookup, rtnpurge, rtnremove
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtnlookup

## look up logical name in table

Displays as a list of hexadecimal numbers the object mapped to a specified logical name in a specified logical name table.

---

**/usr/alliant/rtnlookup** *tblname objname*

---

### ARGUMENTS

*tblname* specifies the name of the table in which to perform the lookup.

*objname* specifies the logical name of the object.

### EXAMPLES

The following example looks up the name x1 in the logical name table xtbl:

```
# /usr/alliant/rtnlookup xtbl x1
Object id: 0x1
```

### DIAGNOSTICS

- privilege violation – You must be super-user. (Rtnlookup accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtnlookup can only be run by super-user.)

### SEE ALSO

rtnamed, rtncrtbl, rtndeltbl, rtndumptbl, rtnenter, rtnremove, rtnpurge
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtnpurge

## purge object from logical name tables

Removes all name-to-object mappings for a specified object from all logical name tables.

---

## /usr/alliant/rtnpurge *objid*

---

### ARGUMENTS

*objid* specifies the object as a list of hexadecimal numbers separated by white space and enclosed in single or double quotes. The quotes are not necessary if only one number is specified.

### NOTES

Neither this command nor the system call name_purge will purge a mapping for an event or semaphore which is marked "permanent."

### EXAMPLES

The following example purges all objects with the identification 0x1, 0x2, 0x3:

```
# /usr/alliant/rtnpurge "1 2 3"
```

### DIAGNOSTICS

- privilege violation – You must be super-user. (Rtnpurge accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtnpurge can only be run by super-user.)

### SEE ALSO

rtnamed, rtncrtbl, rtndeltbl, rtndumptbl, rtnenter, rtnlookup, rtnremove
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rtnremove

## remove logical name from table

Removes a name-to-object mapping from a logical name table.

---

## /usr/alliant/rtnremove *tblname objname*

---

### ARGUMENTS

*tblname* specifies the name of the logical name table containing the name.

*objname* specifies the logical name of the object.

### NOTES

The command rtnremove or the system call name_remove can be used to delete the name-to-object mapping for any object, even if that object is permanent. If the last reference to a permanent event or semaphore is removed from the name table, and no existing process has established a local reference to the object, there is no way of deleting the object without rebooting Concentrix.

### EXAMPLES

The following example removes the object mapped to the logical name x1 in the table xtb:

```
# /usr/alliant/rtnremove xtbl x1
```

### DIAGNOSTICS

- privilege violation – You must be super-user. (Rtnremove accesses the logical name server, rtnamed, via direct rpc calls and does not use the system call interface. For the server to trust the authentication parameters that identify the user, rtnremove can only be run by super-user.)

### SEE ALSO

rtnamed, rtncrtbl, rtndeltbl, rtndumptbl, rtnenter, rtnlookup, rtnpurge
*Concentrix-2800 System Reference: System Calls and C Library Routines*:
name_crtbl, name_deltbl, name_enter, name_lookup, name_purge, name_remove

# rusersd

network username server

Returns a list of users on the network.

The rusersd daemon is started as needed by inetd.

---

## /usr/etc/rpc.rusersd

---

### SEE ALSO

inetd
*Concentrix-2800 Commands and Applications:* rusers
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: services

# rwalld

## network rwall server

Handles rwall and shutdown requests by calling wall for all the appropriate network machines.

The rwalld daemon is started as needed by inetd.

---

## /usr/etc/rpc.rwalld

---

### SEE ALSO

inetd, shutdown
*Concentrix-2800 Commands and Applications Manual:* rwall, wall
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: services

# rwhod

## system status server

Maintains the data base used by the rwho and ruptime programs. Its operation is predicated on the ability to broadcast messages on a network.

The rwhod daemon is typically invoked by /etc/rc.

---

## /etc/rwhod

---

### DESCRIPTION

Rwhod operates as both a producer and consumer of status information. As a producer of information, it periodically queries the state of the system and constructs status messages that are broadcast on a network. As a consumer of information, it listens for other rwhod servers' status messages, validates them, and records them in a collection of files located in the directory /usr/spool/rwho.

Rwho transmits and receives messages at the port indicated by the rwho service specification in /etc/services. The messages sent and received are of the form:

```
struct       outmp {
       char  out_line[8]; /* tty name */
       char  out_name[8]; /* user id */
       long  out_time;    /* time on */
};
struct       whod {
       char  wd_vers;
       char  wd_type;
       char  wd_fill[2];
       int   wd_sendtime;
       int   wd_recvtime;
       char  wd_hostname[32];
       int   wd_loadav[3];
       int   wd_boottime;
       struct       whoent {
              struct       outmp we_utmp;
              int   we_idle;
       } wd_we[1024 / sizeof (struct whoent)];
};
```

All fields are converted to network byte order prior to transmission. The load averages are as calculated by the w program, and represent load averages over the 5, 10, and 15 minute intervals prior to a server's transmission; they are multiplied by 100 for representation in an integer. The host name included is that returned by the gethostname system call, with any trailing domain name omitted. The array at the end of the message contains information about the users logged in to the sending machine. This information includes the contents of the utmp entry for each non-idle terminal line and a value indicating the time in seconds since a character was last received on the terminal line.

Messages received by the rwho server are discarded unless they originated at an rwho server's port. In addition, if the host's name, as specified in the message, contains any unprintable ASCII characters, the message is discarded. Valid messages received by rwhod are placed in files named whod.hostname in the directory /usr/spool/rwho. These files contain only the most recent message, in the format described above.

Status messages are generated approximately once every 3 minutes. Rwhod performs an nlist on /vmunix every 30 minutes to guard against the possibility that this file is not the system image currently operating.

## SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: rwho, ruptime, w
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: nlist
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services, utmp

# sa

## maintain system accounting files

Reports on, cleans up, and maintains system accounting files.

---

**/etc/sa [-abcdDfijkKlmnrstuv] [-S *savacctfile*] [-U *usracctfile*] [*file*]**

---

### ARGUMENTS

**-a** prints all command names, even those containing unprintable characters and those used only once. By default, these are placed under the name ***other.

**-b** sorts output by the sum of user and system time divided by the number of calls. The default sort is by sum of user and system times.

**-c** prints the percentage of total time over all commands, along with total user, system, and real time for each command.

**-d** sorts by average number of disk I/O operations.

**-D** prints and sorts by total number of disk I/O operations.

**-f** suppresses interactive threshold compression with -v flag.

**-i** suppresses reading of a summary file.

**-j** gives seconds per call, instead of total minutes for each category.

**-k** sorts by cpu-time average memory usage.

**-K** prints and sorts by cpu-storage integral.

**-l** separates system and user time; normally they are combined.

**-m** prints the number of processes and number of CPU minutes for each user.

**-n** sorts by number of calls.

**-r** reverses the order of the sort.

**-s** merges the accounting file into the summary file when done.

**-S** *savacctfile*
specifies the summary file; defaults to /usr/adm/savacct.

**-t** reports the ratio of real time to the sum of user and system times for each command.

**-u** supersedes all other flags; prints the user ID and command name for each command in the accounting file.

**-U** specifies the name of the file to accumulate the per-user statistics printed by the -m option; defaults to /usr/adm/usracct.

**-v***n* types the name of each command used *n* times or fewer. Awaits a reply from the terminal; respond y to add the command to the category **junk**. Used to strip out garbage.

*file* specifies an accounting file; /usr/adm/acct is the default.

## DESCRIPTION

Sa condenses the information in /usr/adm/acct into a summary file /usr/adm/savacct containing a count of the number of times each command was called and the time resources were consumed. This condensation is desirable because on a large system /usr/adm/acct can grow by 100 blocks per day. The summary file is normally read before the accounting file, so the reports include all available information.

Output fields are labeled:

| | |
|---|---|
| `cpu` | sum of user+system time (in minutes) |
| `re` | real time (also in minutes) |
| `k` | cpu time averaged core usage (in 1K units) |
| `avio` | average number of I/O operations per execution |
| `tio` | total I/O operations |
| `k*sec` | cpu storage integral (kilo-core seconds) |
| `u` | user cpu time (in minutes) |
| `s` | system cpu time (in minutes) |

If a process is executing on a complex, the base (actual) time is multiplied by the number of processors in the complex. A process which runs for 10 minutes on a complex of 6 CEs accounts for 60 minutes of CPU time.

## FILES

| | |
|---|---|
| /usr/adm/acct | execution accounting file |
| /usr/adm/savacct | summary of execution accounting file |
| /usr/adm/usracct | per-user summary of execution accounting file |

## SEE ALSO

ac, accton
*Concentrix-2800 System Reference: System Calls and C Library Routines*: acct

# sendmail

## send mail over the internet

Sends a message to one or more users, routing the message over whatever networks are necessary. Sendmail does internetwork forwarding as necessary to deliver the message to the correct place. Sendmail is not a user interface program but only delivers pre-formatted messages. Other programs such as mail provide user-friendly interfaces.

---

**/usr/lib/sendmail** [*options*] [*address ...*]

---

### ARGUMENTS

**-ba** enters ARPANET mode. All input lines must end with a carriage return and line feed, and all messages are generated with a carriage return and line feed at the end. Also, the From: and Sender: fields are examined for the name of the sender.

**-bd** runs as a daemon. Sendmail forks and runs in the background listening on socket 25 for incoming SMTP connections. The sendmail daemon is typically invoked by /etc/rc.local.

**-bi** identical to newaliases. (Initializes the alias database.)

**-bm** delivers mail in the usual way (default).

**-bp** identical to mqueue. (Printsa summary of the mail queue.)

**-bs** uses the SMTP protocol as described in RFC821. This option implies all the operations of the -ba option that are compatible with SMTP.

**-bt** runs in address test mode. This mode reads addresses and shows the steps in parsing, and is used for debugging configuration tables.

**-bv** verifies names only; does not try to collect or deliver a message. Verify mode is normally used for validating users or mailing lists.

**-bz** creates the configuration freeze file. **NOTE:** Do not use this option on Alliant systems. The sendmail facility does not work correctly if a frozen configuration file exists.

**-C***file* uses an alternate configuration file. Sendmail refuses to run as root if an alternate configuration file is specified. The frozen configuration file is bypassed.

**-d***X* sets debugging value to *X*.

**-F** *fullname*
sets the full name of the sender.

**-f***name*
sets the name of the sender of the mail. This option is effective only for the user executing sendmail and trusted users (typically, root, daemon, and network).

-h*cnt* sets the hop count to *cnt*. The hop count is incremented each time the mail is processed. When it reaches a limit (by default 30), the mail is returned with an error message.

-M*id* attempts to deliver the queued message with message identification *id*.

-n suppresses aliasing.

-o*xvalue*
sets option *x* to the specified *value*. See "Configuration Options" below. These options can also be set in the configuration file.

-q[*time*]
processes saved messages in the queue at given intervals. If omitted, processes the queue once. The argument is an integer suffixed by one of the following letters: s (seconds), m (minutes), h (hours), d (days), or w (weeks). For example, -q1h30m or -q90m both set the timeout to one hour, thirty minutes.

-r*name*
is an alternate and obsolete form of the -f option.

-R*string*
attempts to deliver from the queue of pending mail any message with a recipient containing the specified string. Useful for clearing out mail directed to a machine that has been down for awhile.

-t reads a message and sends it to the users listed on the To:, Cc:, and Bcc: lines. The Bcc: line is deleted before transmission. Any addresses in the argument list are suppressed.

-v enters verbose mode. Alias expansions are announced, and so forth.

## CONFIGURATION OPTIONS

A*file* specifies an alternate alias file.

c does not initiate immediate connection on mailers that are considered "expensive" to connect to. This option requires queueing.

d*x* sets the delivery mode to *x*. Delivery modes are:

- **i** interactive (synchronous) delivery
- **b** background (asynchronous) delivery
- **q** queue only – that is, actual delivery occurs the next time the queue is run.

D tries to automatically rebuild the alias database if necessary.

**e***x* sets error processing to mode *x*. Valid modes are:

**m** mails back the error message.

**w** writes back the error message (or mails it back if the sender is not logged in).

**p** prints back the error message on the terminal (default).

**q** throws away error messages (only exit status is returned).

**e** mails back error messages with exit status of 0.

If the text of the message is not mailed back by modes **m** or **w** and if the sender is local to the machine, a copy of the message is appended to the file **dead.letter** in the sender's home directory.

**F***mode* specifies in octal the mode to use when creating temporary files, for example, 644 or 600.

**f** saves UNIX-style `From` lines at the front of messages (normally considered redundant).

**g***gid* specifies the default group identification to use when calling mailers.

**H***file* specifies the SMTP help file.

**i** does not interpret periods on lines by themselves as message terminators.

**L***n* specifies the default log level.

**m** sends to **me** (the sender) also in an alias expansion.

**o** allows this message to have old style headers. An adaptive algorithm is used that correctly determines the header format in most cases. Otherwise, this message is guaranteed to have new style headers (that is, commas instead of spaces between addresses).

**Q***queuedir*
specifies the directory in which to queue messages.

**r***timeout*
specifies the timeout on reads; if none is set, **sendmail** waits forever for a mailer.

**S***file* specifies a log file for statistics.

**s** always represents an instance of the queue file, even under circumstances where it is not strictly necessary.

**T***time* specifies the timeout on messages in the queue. A message is returned to the sender if it is still in the queue after the interval passes. The default is three days.

t*stz,dtz* sets the name of the time zone.

u*uid* sets the default user identification for mailers (mailers without an S option in the mailer definition).

## DESCRIPTION

By default, sendmail reads standard input up to an end-of-file or a line with a single dot and sends a copy of the message to all of the addresses listed. The network(s) to use is based on the syntax and contents of the addresses.

Local addresses are looked up in the local aliases file or by using the NIS name service, and aliased appropriately. In addition, if a .forward file is in a recipient's home directory, sendmail forwards a copy of each message to the list of recipients that the file contains. Aliasing can be prevented by preceding the address with a backslash. Normally the sender is not included in any alias expansions.

If the first character of the user name is a vertical bar, the rest of the user name is used as the name of a program to pipe the mail to. Quote the name of the user to keep sendmail from suppressing the blanks from between arguments.

sendmail will also route mail directly to other known hosts in a local network. The list of hosts to which mail is directly sent is maintained in the file /usr/lib/mailhosts.

If invoked as newaliases, sendmail rebuilds the alias database. If invoked as mailq, sendmail prints the contents of the mail queue.

## DIAGNOSTICS

An exit status is returned. The codes are defined in /usr/include/sysexits.h.

- EX_NOHOST – The specified host does not exist.
- EX_NOUSER – The specified user does not exist.
- EX_OK – Successful completion on all addresses.
- EX_OSERR – A operating system error, such as cannot fork or cannot create pipe, occurs.
- EX_SOFTWARE – An internal software error occurs. The error can be a bad argument.
- EX_SYNTAX – A syntax error occurs in an address.
- EX_TEMPFAIL – A message cannot be sent immediately and is queued.
- EX_UNAVAILABLE – A service is unavailable. Also used as a catchall when no other message applies.

## FILES

Except for /etc/sendmail.cf, these pathnames are specified in /etc/sendmail.cf and so could be modified.

| | |
|---|---|
| /etc/aliases | raw data for alias names |
| /etc/aliases.pag | data base of alias names |
| /etc/aliases.dir | binary representation of aliases file |
| /usr/lib/mailhosts | list of hosts to which mail can be sent directly |
| /etc/sendmail.cf | configuration file |
| /etc/sendmail.fc | frozen configuration |
| /etc/sendmail.hf | help file |
| /etc/sendmail.st | collected statistics |
| /usr/bin/uux | to deliver uucp mail |
| /usr/bin/mail | to deliver local mail |
| /usr/spool/mqueue/* | temp files and queued mail |
| ~/.forward | list of recipients for forwarding messages |

## SEE ALSO

mailq, newaliases
*Concentrix-2800 Commands and Applications Manual*: biff, mail, rmail
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: aliases

# setclass

## set process class

Assigns a scheduling class to a process. You must meet at least one of the following criteria: you are super-user; setclassdef all is in effect; the class is a default scheduling class.

---

## /usr/alliant/setclass *class pid*

---

### ARGUMENTS

*class* specifies a scheduling class (an integer in the range 1 to 32) or a resource type (cce, dce, any, or ipo).

*pid* specifies the process identification number.

### DESCRIPTION

By default, assignments of processes to classes take place in the fork and exec system calls. Fork propagates the class of the parent process to the child process. Exec assigns a class to a process as described by the image.

The command execute initiates a process with a specified class through the -cl option.

At login, default scheduling classes are established as follows:

1 (cce) concurrent CE class. FX/C-2800 programs using concurrency and FX/Fortran-2800 programs execute in this class by default (unless the -uniproc compiler option or the -nc loader option is in effect).

2 (dce) detached CE class. No programs generated by Alliant compilers execute in this class by default.

3 (any) generic class. FX/C-2800 programs not using concurrency and Scalar/C programs execute in this class by default. Also, FX/C-2800 and FX/Fortran programs generated with the -uniproc compiler option or the -nc loader option in effect execute in this class.

4(ipo) IP only resource type. Daemons and other system software.

You can reassign default scheduling classes on a per-process basis with the command /usr/alliant/setdefclass.

You can use either the integer value or the mnemonic (cce, dce, any, ipo) to refer to a default scheduling class.

## EXAMPLES

The following example assigns process 8035 to scheduling class 9:

```
# /usr/alliant/setclass 9 8035
```

The following example assigns process 8035 to the default scheduling class for the concurrent CE (cce) resource type, typically class 1:

```
# /usr/alliant/setclass cce 8035
```

## DIAGNOSTICS

- `Invalid argument` – The class number is invalid.
- `Not owner` – You must be superuser to use **setclass** unless the target scheduling class is a default class or **setclassdef all** is in effect.
- `No such process` – The process identification number is invalid.

## SEE ALSO

setcomplex, setdefclass, setsched, showdefclass, showsched
*Concentrix-2800 Commands and Applications*: execute, login
*Concentrix-2800 System Reference: System Calls and C Library Routines*: execven, getclass, getdefclass, setclass, setdefclass, setclassdef, setsched

# setclassdef

## change setclassdef flag in kernel image

Sets the value of the setclassdef flag; this flag determines whether the changing of job classes (other than the default classes) is restricted to superusers or is open to all users.

---

**/usr/alliant/setclassdef** [*flag* [*kernel-name*]]

---

### ARGUMENTS

*flag* One of two values:

**su** (default) only the superuser may change job classes.

**all** all users may change job classes.

*kernel-name*
(defaults to /unix) name of the kernel image.

### EXAMPLES

The following example permits all users to change job classes:

```
# /usr/alliant/setclassdef all
```

### DIAGNOSTICS

- `The file` *filename* `does not exist`

### NOTES

- You must rerun setclassdef whenever you install a new kernel.
- The script only changes the /unix kernel on disk, so the change is not immediate. You must reboot in order to effect the change.
- To change the setclassdef flag with immediate efect without rebooting, use adb as follows (the example enables all users to use setclass):

```
# adb -w /unix /dev/kmem
ready
setclassdef/W0
_setclassdef:  0x0          =  0x1

#
```

## SEE ALSO

setclass, setsched, showsched
*Concentrix-2800 Commands and Applications:* execute
*Concentrix-2800 System Reference; System Calls and C Library Routiness:* setclass, setsched

# setcomplex

## set or display thresholds for computational complex

Sets or displays the thresholds that determine when a computational complex (CL resource) dynamically (1) explodes into detached CEs and (2) coalesces into concurrent CEs. Specify -c and/or -d to set thresholds. Specify no arguments to display all thresholds. Specify only a CL number to display thresholds for one specified complex. Applies only to the time-sharing scheduler.

---

**/usr/alliant/setcomplex [[CL *number*] [-c*n*] [-d*n*]]**

---

### ARGUMENTS

**CL *number***
specifies a computational complex; *number* must be in the range 0–5.

**-c*n*** (concurrent threshold) specifies the threshold for unserviced concurrent jobs. During a -tc time slice, the number of unserviced concurrent jobs must be less than *n* for the complex to explode. During a -td time slice, the number of unserviced concurrent jobs must be greater than or equal to *n* for the complex to coalesce.

**-d*n*** (detach threshold) specifies the threshold for unserviced generic jobs. During a -tc time slice, the number of unserviced generic jobs must be greater than *n* for the complex to explode. During a -td time slice, the number of unserviced generic jobs must be less than or equal to *n* for the complex to coalesce.

### DESCRIPTION

A computational complex (CL resource) operates in one of two states during each time slice:

- **Coalesced** – In the coalesced state, the CEs that make up the complex work as one entity and are said to be attached.
- **Exploded** – In the exploded state, the CEs that make up the complex work as separate entities and are said to be detached.

Whether the complex coalesces or explodes during a time slice is regulated by the unserviced job load, the -c and -d options of this command, and the -t, -tc, and -td specifications to the CL option of the setsched command. For purposes of coalescing and exploding the complex, the unserviced job load consists of two parts:

- **Concurrent jobs** – Programs that contain concurrency instructions and therefore profit from execution on the concurrent CEs of a coalesced complex.
- **Generic jobs** – Programs that do not contain concurrency instructions and therefore can run as efficiently on a detached CE in an exploded complex.

The algorithm for coalescing and exploding the complex is stated in terms of the CL option of the command /usr/alliant/setsched (-t for neutrality, -tc to favor concurrent work, and -td to favor generic (detached) work); the unserviced job load (number of waiting concurrent jobs versus waiting generic jobs); and the -c*n* (concurrent threshold) and -d*n* (detach threshold) options of this command:

| CL | unserviced job load | complex |
|---|---|---|
| -t | con > gen | coalesces |
| | gen > con | explodes |
| | gen = con | remains as is |
| -tc | gen <= d*n* | coalesces |
| | con >= c*n* | coalesces |
| | gen > d*n* & con < c*n* | explodes |
| -td | con < c*n* | explodes |
| | gen > d*n* | explodes |
| | gen <= d*n* & con >= c*n* | coalesces |

## EXAMPLE

The following example displays the default settings, then makes a change:

```
# /usr/alliant/setcomplex
/usr/alliant/setcomplex CL 0 -d1 -c1
/usr/alliant/setcomplex CL 1 -d1 -c1
# /usr/alliant/setcomplex CL 0 -d2
```

## SEE ALSO

setclass, setsched, showsched
*Concentrix-2800 System Reference: System Calls and C Library Routines*: setclass, setsched

# setcres

## set computing resource

Assigns a computing resource type, a maximum number of processors, and a scheduling class to a specified process. If you are not super-user, you must own the specified process and at least one of the following criteria must be true: setclassdef all is in effect; the class is a default scheduling class.

---

**/usr/alliant/setcres** *pid crestype maxpn class*

---

### ARGUMENTS

*pid* is the process identification number.

*crestype* specifies a computing resource type (cce, dce, any, or ipo).

*maxpn* specifies the maximum number of processors (if concurrent). The value -1 assigns the largest complex size available.

*class* specifies a scheduling class (an integer in the range 1 to 32) or a resource type (cce, dce, any, or ipo). The value -1 uses the default class for the specified computing resource type.

### DESCRIPTION

The setcres command allows a user to change certain characteristics of a process.

**Warning:** The process hangs if it is changed to a resource type or class that is not currently available. Check the scheduling and cluster information with the commands showsched and clinfo before making any changes.

### DIAGNOSTICS

- `Exec format error` - The class number is invalid.
- `Not owner` - You must be superuser to use setclass unless the target scheduling class is a default class or setclassdef all is in effect.
- `No such process` - The process identification number is invalid.
- `Invalid computing resource type` - The computing resource type is not one of: cce, dce, any, or ipo.

## EXAMPLE

The following example assigns process 17516 to run on a concurrent CE resource type, using four processors and the scheduling class 5.

```
# /usr/alliant/setcres 17516 cce  4  5
```

## SEE ALSO

setclass, setpn, setsched, settype,
*Concentrix-2800 System Reference: System Calls and C Routines:* setcres

# setdefclass

## set default scheduling class

Assigns default scheduling classes to resource types on a per-process basis.

---

**/usr/alliant/setdefclass** *pid rtype class* [*rtype class*] ...

---

### ARGUMENTS

*pid* specifies the process identification number.

*rtype* specifies a resource type (cce, dce, any, or ipo). Each resource type has exactly one default scheduling class; if you specify a resource type more than once, the last specification applies.

*class* specifies a scheduling class (an integer in the range 1–32).

### DESCRIPTION

At login, default scheduling classes are established as follows:

1 (cce) concurrent CE class. FX/C-2800 programs using concurrency and FX/Fortran-2800 programs execute in this class by default (unless the -uniproc compiler option or the -nc loader option is in effect).

2 (dce) detached CE class. No programs generated by Alliant compilers execute in this class by default.

3 (any) generic class. FX/C-2800 programs not using concurrency and Scalar/C programs execute in this class by default. Also, FX/C-2800 and FX/Fortran programs generated with the -uniproc compiler option or the -nc loader option in effect execute in this class.

4 (ipo) IP only class. Daemons and other system software execute in this class by default.

### EXAMPLES

The following example assigns class 9 as the default scheduling class for the IP only (ipo) resource type:

```
# /usr/alliant/setdefclass 8035 ipo 9
```

## DIAGNOSTICS

- `Not owner` – You must be superuser to use setdefclass.
- `No such process` – The process identification number is invalid.
- `Invalid argument` – The class number is invalid.

## SEE ALSO

setclass, setcomplex, setsched, showdefclass, showsched
*Concentrix-2800 Commands and Applications*: execute, login
*Concentrix-2800 System Reference: System Calls and C Library Routines*: execve, getclass, getdefclass, setclass, setdefclass, setsched

# setmaxshare

## change maxsharesize flag in kernel image

Sets the value of the maxsharesize kernel variable; this kernel variable determines the maximum amount of physical memory that can be used for shared regions. A specification of zero, which is the default, causes maxsharesize to be calculated as follows when the system boots: the value of usermem (user memory) minus the minimum of (a) 25% of usermem or (b) 12MB.

The new value does not take effect until you reboot the system. Whenever you install a new kernel, you must rerun the setmaxshare command.

---

**/usr/alliant/setmaxshare** [[*value*] [*kernel-name*]]

---

### ARGUMENTS

If you specify no arguments, setmaxshare prints the current value of maxsharesize. If the current object file value is zero, the core value calculated by the system is printed.

*value* is a hexadecimal number (with or without the prefix 0x) that specifies the amount of memory available for shared memory regions in units of 4096 (0x1000) bytes. Do not exceed the value of usermem. A specification of zero specifies the default (calculated by the system as explained above).

*kernel-name*
name of the kernel image; defaults to /unix.

### NOTES ON SHARED REGION LIMITS

Two additional kernel variables affect shared region limits:

- MAX_SHARED_CREATE – specifies the maximum size of a temporary shared region in page-size (4096 bytes) units. By default, this variable is set to 2,097,152 (0x200000) bytes.

- MAX_SHARED_PERMCREATE – specifies the maximum size of a permanent shared region in page-size (4096 bytes) units for ordinary users; the super-user can create pemanent regions of any size. By default, this variable is set to zero, which prohibits ordinary users from creating shared regions.

Use the Patch_kernel procedure to view and adjust these variables. Both variables are dynamic – the system does *not* have to be rebooted to make them take effect.

## EXAMPLE

The following example sets the maximum amount of shared memory regions at 33,554,432 (0x2000000) bytes:

```
# /usr/alliant/setmaxshare 2000
._maxsharesize: 0x0             =       0x2000
/unix image changed.  Reboot the system with this image to enable change.
```

The following example sets the maximum size of a shared memory region at 2,097,152 (0x200000) bytes:

```
# adb -w /unix /dev/kmem
ready
MAX_SHARED_PERMCREATE/W 0x200
._MAX_SHARED_PERMCREATE:        0x0             =       0x200
$q
```

## DIAGNOSTICS

- The file *filename* does not exist
- *filename* image changed. Reboot the system with this image to enable change.

## SEE ALSO

srcreate, srdelete, srdisplay
*Concentrix-2800 System Reference*: *System Calls and C Library Routines:* create_shared_region

# setpn

## set the maximum number of processors for a concurrent job

Sets the maximum number of processors, and optionally a scheduling class to a specified process. If you are not super-user, you must own the specified process and at least one of the following criteria must be met: setclassdef all is in effect; the class is a default scheduling class.

---

**/usr/alliant/setpn** *pid maxpn* [*class*]

---

### ARGUMENTS

*pid* specifies the process identification number.

*maxpn* specifies the maximum number of processors (if concurrent). The value –1 indicates the largest complex size available.

*class* specifies a scheduling class (an integer in the range 1 to 32) or a resource type (cce, dce, any, or ipo). The value –1 uses the default class for the specified resource type.

### DESCRIPTION

The setpn utility allows a user to set the maximum number of processors for a concurrent process and (optionally) its scheduling class. The computing resource type is automatically set to be concurrent.

**Warning:** The process hangs if it is changed to a resource type or class that is not currently available. Check the scheduling and cluster information with showsched and clinfo before using this command.

### EXAMPLE

The following example assigns process 17516 to run on 4 processors in class 5.

```
# /usr/alliant/setpn 17516 4 5
```

### DIAGNOSTICS

- `Exec format error` – The class number is invalid.
- `Not owner` – You must be super-user to use setclass unless the target scheduling class is a default class or setclassdef all is in effect.
- `No such process` – The process identification number is invalid.

### SEE ALSO

setclass, setcres, setsched, settype
*Concentrix-2800 System Reference: System Calls and C Routines:* setcres

# setsched

## set scheduling vector for computing resource

Modifies a computing resource scheduling vector. A scheduling vector determines the classes of work that can be processed by a computing resource. Each resource in the system configuration has its own scheduling vector.

---

**/usr/alliant/setsched** *resource number vector ...*

---

### ARGUMENTS

*resource* specifies the type of resource:

**IP** interactive processor; *number* must be in the range 0–7.

**CE** computational element; *number* must be in the range 0–27.

**CL** computational complex; *number* must be in the range 0–5.

*number* specifies the particular IP, CE, or CL.

*vector* (maximum of 8) specifies a time slice over which the resource is scheduled in one of the following forms:

**-t** *ticks class ...*

**-tc** *ticks class ...* [CL only]

**-td** *ticks class ...* [CL only]

**-tf** *master-frame-scheduler* [FX/RT-2800 only]

**-tfm** *master-frame-scheduler* [FX/RT-2800 only]

**-tr** *ticks class ...* [FX/RT-2800 only]

*ticks* an integer in the range 3–120 that specifies a time slice in increments of 10 milliseconds. A specification less than 3 is silently adjusted to 3; a specification greater than 120 is silently adjusted to 120.

*class* an integer in the range 1–32 that specifies a job class; mnemonics are also available – see below. The list *class ...* specifies up to 32 classes, in order of priority, that can run for the specified number of *ticks*.

*master-frame-scheduler*
an integer that specifies a master frame scheduler.

**-t** uses the time-sharing scheduler. If specified for a computational complex (CL resource), the complex explodes or coalesces based on the unserviced job load. The complex explodes if the number of unserviced generic jobs is greater than the number of unserviced concurrent jobs. The complex coalesces if the number of unserviced concurrent jobs is greater than the number of unserviced generic jobs. Otherwise, the state of the complex does not change.

**-tc** uses the time-sharing scheduler. Causes a computational complex (CL resource) to favor concurrent work; the complex explodes only when the number of unserviced generic jobs is greater than the resource's detach threshold and the number of unserviced concurrent jobs is less than the resource's concurrent threshold. Set thresholds with the command /usr/alliant/setcomplex.

**-td** uses the time-sharing scheduler. Causes a computational complex (CL resource) to favor generic (detached) work; the complex coalesces only when the number of unserviced generic jobs is less than or equal to the resource's detach threshold and the number of unserviced concurrent jobs is greater than or equal to the resource's concurrent threshold. Set thresholds with the command /usr/alliant/setcomplex.

**-tf** assigns this resource to run a real-time frame scheduler under a specified real-time master frame scheduler. You cannot specify the real-time frame scheduler for IP0. You cannot specify the real-time frame scheduler for a CE that is configured as part of a complex. (Returns the error, "File exists.") You can specify the real-time frame scheduler for a statically detached CE and for a complex (CL). This specification must be the only *vector* specification, and entirely devotes the resource to the frame scheduler.

**-tfm** assigns this resource to run a specified real-time master frame scheduler. You cannot specify the real-time frame scheduler for IP0. You cannot specify the real-time frame scheduler for a CE that is configured as part of a complex. (Returns the error, "File exists.") You can specify the real-time frame scheduler for a statically detached CE and for a complex (CL). This specification must be the only *vector* specification, and entirely devotes the resource to the frame scheduler.

**-tr** uses the real-time scheduler. You cannot specify the real-time scheduler for IP0. You cannot specify the real-time scheduler for a CE that is configured as part of a complex. (Returns the error, "File exists.") You can specify the real-time scheduler for a statically detached CE and for a complex (CL).

## DESCRIPTION (-t, -tc, -td, -tr)

The scheduling vector for each process contains a list of time slices each of which is associated with a list of classes. During a time slice, the processor processes the job of highest priority of the first class in the list. If a job of higher class priority within the current time duration becomes available, it can preempt and displace the current job. If no work is available in the highest-priority class, the resource accepts work from the next highest priority class, and so forth. (Real-time classes are not prioritized by order of specification, but by current *real-time* priority.)

Typically, you use -t time slices for time-sharing work on CE, CL, and IP resources; -tr time slices for real-time work on CE and IP resources; and -tc and -td time slices for time-sharing work on CL resources. You can force CL resources into permanent exploded or coalesced states by setting the detach and concurrent thresholds very high; see the command /usr/alliant/setcomplex.

The default scheduling classes used by Alliant programs are as follows. The parenthesized mnemonics can be used in place of the numbers.

1 (cce) concurrent CE class. FX/C-2800 programs using concurrency and FX/Fortran-2800 programs execute in this class by default (unless the -uniproc compiler option or the -nc loader option is in effect).

2 (dce) detached CE class. No programs generated by Alliant compilers execute in this class by default.

3 (any) generic class. FX/C-2800 programs not using concurrency and Scalar/C programs execute in this class by default. Also, FX/C-2800 and FX/Fortran programs generated with the -uniproc compiler option or the -nc loader option in effect execute in this class.

4 (ipo) IP only class. Daemons and other system software execute in this class by default.

## DESCRIPTION (-tf, -tfm)

Before executing a program that starts a frame schedule, ensure that resources are dedicated for the specified master frame schedulers (-tfm) and frame schedulers (-tf). Both the -tfm and -tf options specify the number of the master frame scheduler. Typically, one resource executes the master frame scheduler, and one or more resources execute frame schedulers under that master frame scheduler. The program using the frame schedule specifies the master frame scheduler in the frm_create_sched system call.

The program using the frame schedule also specifies the resource(s) used to run frame schedulers under the master frame scheduler. These resources must in fact be assigned to frame schedulers under the correct master frame scheduler before the program executes.

A program that attempts to execute under a frame schedule without prior assignment of the appropriate resources will hang, and cannot be killed or continued. You must reboot to eliminate the process.

## EXAMPLES

The following example executes a file that contains contains setsched commands permitting the execution of work in classes 9 and 16:

```
# cat /usr/adm/sched.alt
setsched IP  0 -t  100 4
setsched IP  1 -t  40 9 4 3 16 -t 20 9 3 4 16
setsched CL  0 -tc 30 9 1 2 3 16 -td 30 9 2 3 1 16 -td 30 9 3 2 1 16
setsched CE  0 -t  30 9 2 3 16 -t 30 9 3 2 16
setsched CE  1 -t  30 9 2 3 16 -t 30 9 3 2 16
setsched CE  2 -t  30 9 2 3 16 -t 30 9 3 2 16
setsched CE  3 -t  30 9 2 3 16 -t 30 9 3 2 16
# chmod 744 /usr/adm/sched.alt
# cd /usr/alliant
# /usr/adm/sched.alt
```

The following example dedicates CE3 to run real-time master scheduler #1, and dedicates CE0 to run a real-time scheduler under master scheduler #1:

```
# setsched CE 0 -tf 1
# setsched CE 3 -tfm 1
```

## DIAGNOSTICS

- `Permission denied` – You must be licensed for FX/RT-2800 to use the –tr, –tf, or –tfm option.

## SEE ALSO

setclass, setcomplex, setsched, showdefclass, showsched
*Concentrix-2800 Commands and Applications*: execute, login
*Concentrix-2800 System Reference: System Calls and C Library Routines*: execven, getclass, getdefclass, setclass, setdefclass, setsched

# settype

## set computing resource type

Assigns a computing resource type and optionally a scheduling class to a specified process. If you are not super-user, you must own the specified process and at least one of the following criteria must be met: setclassdef all is in effect; the class is a default scheduling class.

---

**usr/alliant/settype** *pid crestype* [*class*]

---

### ARGUMENTS

*pid* specifies the process identification number.

*crestype* specifies a computing resource type (cce, dce, any, or ipo).

*class* specifies a scheduling class (an integer in the range 1 to 32) or a resource type (cce, dce, any, or ipo). The value -1 (or omitting this argument) assigns the default class for the specified computing resource type.

### DESCRIPTION

The settype command allows a user to change the computing resource type and (optionally) the class of a process.

**Warning:** The process hangs if it is changed to a resource type or class that is not currently available. Check the scheduling and cluster information with the showsched and clinfo commands before using this command.

### DIAGNOSTICS

- `Exec format error` - The class number is invalid.
- `Not owner` - You must be super-user to use setclass unless the target scheduling class is a default class or setclassdef all is in effect.
- `No such process` - The process identification number is invalid.
- `Invalid computing resource type` - The computing resource type is not one of: cce, dce, any, or ipo.

### SEE ALSO

setclass, setcres, setpn, setsched
*Concentrix-2800 System Reference: System Calls and C Routines:* setcres

# setup_client

## add or remove diskless client to/from server

Installs or removes a diskless client's environment on the server. An installation includes making necessary directories, updating configuration files, and actually installing the client's file space on the server. The items are removed if the client is being deleted.

---

**/usr/etc/install/setup_client** *op name yptype swapsize rootpath swappath dumppath homepath execpath sharepath arch*

---

### ARGUMENTS

Running the command with no arguments lists the available options.

*op* operation type: add or remove.

*name* *hostname* of the client being added or removed.

*yptype* NIS mode of operation for the client: master, slave, client, or none.

*swapsize* size of the swap file to make for this client. This value is specified in bytes by default, or the letters M (x 1048576), K (x 1000), or B (x 512) may be appended to the value to multiply by the value given.

*rootpath, swappath, dumppath*
full pathnames to be used for each of these functions (client root file system, swap file for client, and file to use for crash dumps). The word none can be specified for *dumppath*, in which case the client kernel uses *swappath* when generating crash dumps. *Rootpath* and *swappath* can also be specified as none for the remove operation. Any pathname so entered will not be removed during the operation.

*homepath* location of user home directories and is typically /home. If another server is providing this service, the pathname can have the server name and colon prepended, as in *servername*:/home. Can also be specified as none if this feature is not used.

*execpath* full pathname of the location on the server where the /usr filesystem is located for this client's architecture type, for example, /export/execs/sun3.

*sharepath* full pathname of the location on the server where the shared /usr/share filesystem is located, for example, /export/exec/share or none. Setup_exec takes no action with this value, but passes it along to create_root.

*arch* architecture name for this client, for example, sun2, sun3, and sun4.

### DIAGNOSTICS

Error output is self-explanatory.

## FILES

/tftpboot/*

## SEE ALSO

create_root, fix_bootparams, INSTALL, sunboot
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats:* exports

# setup_exec

## install diskless client software on a server

Loads software for a particular diskless client architecture from the manufacturer's software distribution tapes. Typically called by INSTALL, but can be invoked manually to install a particular architecture type.

---

**setup_exec** *arch execpath tapetype* [*tapehost*] *tapedev sharepath*

---

### ARGUMENTS

Running the command with no arguments lists the available options.

*arch* architecture name to be loaded, for example, sun2, sun3, and sun4.

*execpath* full pathname of the location on the server to load the software, for example, /exports/execs/sun3.

*tapetype* if remote, the *tapehost* argument is taken to be the hostname of the machine to which the tape drive is attached. This hostname is verified via NIS (if domainname is not NULL or set to noname) or /etc/hosts.

*tapedev* tape driver type being used (such as mt, st, or ar), or else the full pathname of the non-rewinding raw tape drive.

*sharepath* full pathname of the location on the server to load that part of the software that can be shared among different architectures (typically the /usr/share tree), for example, /export/exec/share. Shared software can be disabled by specifying a value of none in this field.

### NOTES

Some tape devices have a bug that randomly causes unexpected results during a seek to a tape file if the tape is not positioned at beginning of tape. To handle this case reliably requires that the tape be rewound after every file read and before seeking to the next file to be read, which is much slower than attempting to perform seeks from where the tape was last left.

The default action of setup_exec is to perform a rewind between every tape file read. An internal variable at the beginning of the script called DOVERIFY can be changed to take advantage of mid-tape seeks. See the comments in the script to take advantage of this feature, but change it at your own risk.

### DIAGNOSTICS

Error output is self-explanatory.

### FILES

/tmp/*.<architecture>

### SEE ALSO

extracting, INSTALL, opt_software, verify_tapevol_arch

# showdefclass

## display default scheduling classes

Displays the default scheduling classes for a process in the form of a setdefclass command.

---

## /usr/alliant/showdefclass *pid*

---

### ARGUMENTS

*pid* specifies the process identification number.

### DESCRIPTION

At login, default scheduling classes are established as follows:

1 (cce) concurrent CE class. FX/C-2800 programs using concurrency and FX/Fortran-2800 programs execute in this class by default (unless the -uniproc compiler option or the -nc loader option is in effect).

2 (dce) detached CE class. No programs generated by Alliant compilers execute in this class by default.

3 (any) generic class. FX/C-2800 programs not using concurrency and Scalar/C programs execute in this class by default. Also, FX/C-2800 and FX/Fortran programs generated with the -uniproc compiler option or the -nc loader option in effect execute in this class.

4 (ipo) IP only class. Daemons and other system software execute in this class by default.

You can reassign default scheduling classes on a per-process basis with the command /usr/alliant/setdefclass.

### EXAMPLES

The following example displays the default scheduling classes for process 8035:

```
# /usr/alliant/showdefclass 8035
setdefclass 8035 cce 1 dce 2 any 3 ipo 4
```

## DIAGNOSTICS

- `Not owner` – You must be superuser to use showdefclass.
- `No such process` – The process identification number is invalid.

## SEE ALSO

setclass, setcomplex, setdefclass, setsched, showsched
*Concentrix-2800 Commands and Applications*: execute, login
*Concentrix-2800 System Reference: System Calls and C Library Routines*: execven, getclass, getdefclass, setclass, setdefclass, setsched

# showmount

## show all remote mounts

Lists all clients that remotely mounted a filesystem from a specified host. This information is maintained by the mountd server on *host* and is saved across crashes in the file /etc/rmtab.

---

**/etc/showmount [-a] [-d] [-e] [*host*]**

---

### ARGUMENTS

**-a** prints all remote mounts in the format *hostname:directory*, where *hostname* is the name of the client and *directory* is the root of the file system that was mounted.

**-d** lists directories remotely mounted by clients.

**-e** prints the list of exported file systems.

*host* specifies the name of a host. The default is the value returned by the hostname command.

### NOTES

- If a client crashes, its entry is not removed from the list until it reboots and executes umount -a.

### FILES

/etc/rmtab

### SEE ALSO

mountd
*Concentrix-2800 Commands and Applications Manual:* hostname
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: rmtab, exports

# showsched

## display scheduling vectors

Displays the current values of the system's computing resource scheduling vectors in the form of setsched commands. You can use redirection to capture showsched output in a file that you can then edit and execute as a script (change mode to allow execution).

---

## /usr/alliant/showsched [-m]

---

### ARGUMENTS

**-m** uses the mnemonics cce, dce, any, and ipo in place of the numeric values for the default classes.

### EXAMPLE

The following example demonstrates the default settings of the scheduling vectors for a system with two IPs, eight CEs, and one cluster (CL):

```
# /usr/alliant/showsched
setsched IP  0 -t  60  4
setsched IP  1 -t  40  4  3
setsched CL  0 -tc 30  1  2  3 -td 30  1 3 2
setsched CE  0 -t  30  2  3 -t  30  3  2
setsched CE  1 -t  30  2  3 -t  30  3  2
setsched CE  2 -t  30  2  3 -t  30  3  2
setsched CE  3 -t  30  2  3 -t  30  3  2
setsched CE  4 -t  30  2  3 -t  30  3  2
setsched CE  5 -t  30  2  3 -t  30  3  2
setsched CE  6 -t  30  2  3 -t  30  3  2
setsched CE  7 -t  30  2  3 -t  30  3  2
```

### SEE ALSO

setclass, setsched
*Concentrix-2800 System Reference: System Calls and C Library Routines*: setclass, setsched

# shutdown

## shut down system at a given time

Shuts down the system after optionally giving users advance notice, or performs other actions depending on the arguments. If no arguments are specified, the system shuts down to single-user mode in Concentrix-2800. Type sync[RETURN], /etc/umount -a[RETURN], sync[RETURN], and [CTRL]d to exit from single-user Concentrix-2800 and bring up multi-user Concentrix-2800. Logs the shutdown in the system log unless -n is specified.

---

**/etc/shutdown** [*argument*] *time* [*message ...*]

---

### ARGUMENTS

**-B** stops in VMEbus booter mode and prompts for names for the Diagnostix-2800 kernel, the Diagnostix-2800 kernel debugger, and the Diagnostix-2800 root file system. Type [RETURN] to continue with default file names. You can skip loading the Diagnostix-2800 debugger by typing [ESCAPE]. From the booter prompt, you can also type %h to get help and %ls to list the root directory. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-b** stops in FX/2800 booter mode and prompts for file names for the Concentrix-2800 kernel, the Concentrix-2800 kernel debugger, and the processor configuration file. Type [RETURN] to continue with default file names. From the booter prompt, you can also type ! (exclamation mark) to connect to Diagnostix-2800 (type [CTRL]d to return to the FX/2800 booter). This option causes Concentrix-2800 to be rebooted.

**-D** stops in Alliant ddt mode and waits for commands. You can proceed by typing [ESCAPE]p or [ESCAPE]g. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-d** stops in Alliant ddt 2800 mode and waits for commands. You can proceed by typing [ESCAPE]p or [ESCAPE]g. This option causes Concentrix-2800 to be rebooted.

**-f** creates a file /fastboot so that the system skips file system checking on the next reboot.

**-H** halts the system (shuts the system down to console mode). Recommended for power-down. From the console, type u to exit from console mode and start the system. From a remote terminal, type o r[RETURN]u.

**-h** shuts the system down to Diagnostix-2800 mode. To start Concentrix-2800, type [CTRL]d, then respond y to the FX/2800 prompt.

**-k** simulates but does not actually execute a shutdown.

**-n** does not sync the disks (for emergencies only); inhibits the log message.

**-r** reboots both Diagnostix-2800 and Concentrix-2800.

**-S** stops in Diagnostix-2800 mode. To continue, type CTRL d, then respond y to the FX/2800 prompt. This option causes both Diagnostix-2800 and Concentrix-2800 to be rebooted.

**-s** stops at single-user mode in Concentrix-2800. To continue, type sync RETURN, /etc/umount -a RETURN, sync RETURN, and CTRL d. This option causes Concentrix-2800 to be rebooted.

*time* specifies the shutdown time in one of the following formats:

*+number* a number of minutes from the current time, where *number* is an integer.

*hour:min* at the specified hour (24-hour clock) and minute.

**now** immediately.

*message*
specifies a message, usually the reason for shutting down and the length of the shutdown.

## DESCRIPTION

At intervals that get closer as shutdown approaches, warnings are broadcast to the terminals of all users on the system. A warning consists of the shutdown time and the specified message (if any). Shutdown causes the following actions:

- **Nologin** – Five minutes before shutdown, or immediately if shutdown is in less than five minutes, a warning is written to the file /etc/nologin. If this file exists when a user attempts to log in, login displays its contents and exits. Shutdown removes the file just before exiting.

- **Shutdownlog** – At shutdown time a message is written in the system log, containing the time, the notation Shutdown:, the message, and the name of the user who invoked shutdown.

## NOTES

- You can only shut down the system between now and 23:59 if you specify an absolute time.

## EXAMPLES

The following example shuts down the system in five minutes and immediately reboots. Interactive users receive the message to reconfigure – immediate reboot as part of the shutdown warning.

```
# /etc/shutdown -r +5 to reconfigure - immediate reboot
```

The following example shuts down the system at 1 P.M. and immediately reboots.

```
# /etc/shutdown -r 13:00 to reconfigure - immediate reboot
```

The following example immediately shuts down the system to single-user mode.

```
# /etc/shutdown now
```

The following example immediately halts the system (leaving you in console mode):

```
# /etc/shutdown -H now
```

## FILES

| | |
|---|---|
| /dunix | Diagnostix-2800 kernel image |
| /etc/nologin | file whose existence prevents logins |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /vmunix | a symbolic link to /unix |

## SEE ALSO

fastboot, fsck, init, rc, reboot
*Concentrix-2800 Commands and Applications Manual*: login

# software-install

## install layered product

Installs a layered product in a directory tree starting at *directory/product-name*.

---

**/usr/alliant/software-install** [*product-name* [*directory* [*device*]]]

---

### ARGUMENTS

*product-name*
: specifies the name of the layered productm, or the keyword all. Defaults to all.

*directory*
: specifies the absolute pathname of the directory under which the product is to be installed. The directory must exist and must exist in a non-root file system. Defaults to /usr/installs.

*device*
: specifies the device that contains the distribution volume. Defaults to /dev/rschs.

### DESCRIPTION

Best practice is to follow the installation guide for the product you are installing. The generic procedure for installing any product is listed below:

1. Log in as `root`.
2. (optional but recommended) Go to single-user Concentrix-2800 by typing, for example:

   ```
   # /etc/shutdown +30 to install layered product
   ```

3. If a version of the product is already installed on your system, you may want to archive or move the directory tree starting at *directory/product-name*.
4. Load the distribution tape. Set the write-protection switch on the 8mm Exabyte cartridge so that the red tab is in place, insert the cartridge, and close the door.
5. Invoke software-install.

   If the specified directory does not exist, you can have software-install create the directory or abort. If the specified directory exists but is not a directory, software-install aborts.

   If the specified device does not exist, software-install aborts.

6. The following messages appear:

   ```
   The destination directory selected for installation is: directory
   The device selected to read the distribution media is: device
   The media should be mounted and online if necessary.
   Do you wish to continue? (y/n):
   ```

   If the names are correct, ready *device* (if it is not ready), and type y followed by RETURN. Otherwise, type n followed by RETURN and restart at step 5.

7. The following message appears:

```
.... Checking for old version(s) of Layered Software.....
```

Possible error messages:

```
Directory for Layered Software Product product-name already exists.
This usually means that an old version of software was previously
installed.
type: `r' to remove old version and directory
      `o' to overwrite existing version
      `q' to quit installation
```

If you have archived or copied the old version, type r or o followed by RETURN and follow any additional instructions. Otherwise, type q followed by RETURN and restart at step 3. If you select r and the old installation script does not exist, you can elect to continue or abort.

8. The installation proceeds with the following message:

```
.... Restoring Layered Software product(s): product-name ...
```

When the following message appears, you can remove the distribution volume (store it in a safe place):

```
.... You may now unmount the media from device.
```

If the distribution volume does not contain the expected installation script, software-install aborts.

9. The installation proper proceeds with the following message. The actions and duration of the procedure depend on the product being installed.

```
.... Executing the installation script for product-name.
```

When the installation proper completes, you are afforded an opportunity to log the installation in /usr/installs/install.log:

```
Do you wish to create a new log entry for product-name? (y/n):
```

Type y to log the installation or n to skip the log entry. A log entry contains the following information: date, product name, release, revision, installation device, installer, and installation directory. If the log file does not exist (and you type y), the file /usr/installs/install.log is first created and initialized. If the release number and revision letter of the product are not listed in the installation script, you are queried for this information. You are also queried for your name.

At the end of a successful product installation, the following message appears:

```
.... Installation of product-name complete.
```

This step repeats for each product being installed.

10. When all products are installed, the script terminates with the following messages:

```
.... Software-install procedure complete.
.... Exiting software-install procedure.
```

As with all software installations, you should back up the partition containing the newly installed product immediately.

## EXAMPLE

The following example installs FX/C-2800 in the directory tree starting at /user/fxc-2800 from the default tape:

```
# software-install fxc-2800 /user
```

# spray

## spray packets

Sends a one-way stream of packets to *host* using RPC, and then reports the number received by *host* and the transfer rate.

---

**/usr/etc/spray** *host* [**-c** *count*] [**-d** *delay*] [**-i** *delay*] [**-l** *length*]

---

### ARGUMENTS

*host* a name or an Internet address.

-c *count* specifies how many packets to send; defaults to the number of packets required to make the total stream size 100,000 bytes.

-d *delay* specifies how may microseconds to pause between sending each packet; defaults to 0.

-i uses ICMP echo packets rather than RPC. Since ICMP automatically echos, this option creates a two-way stream.

-l *length* specifies the numbers of bytes in the Ethernet packet that holds the RPC call message; defaults to 86 (the size of the RPC and UDP headers). Spray is encoded using XDR, which only deals with 32-bit quantities, and so rounds up to the nearest possible value. When the length is greater than 1514, the RPC call can not be encapsulated in one Ethernet packet, and the length field no longer has a simple correspondence to Ethernet packet size.

### SEE ALSO

ping, sprayd
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: icmp

# sprayd

## spray server

Records the packets sent by spray.

The sprayd daemon is started as needed by inetd.

---

**/usr/etc/rpc.sprayd**

---

### SEE ALSO

inetd, spray

# srboot

## Alliant FX/2800 booter

Loads and starts Concentrix-2800.

**Note:** This program runs under Diagnostix-2800.

---

**/console/srboot** *unix* [*ddt*] [**-c** *cfg*] [*option ...*]

---

### ARGUMENTS

*unix* specifies the name of the Concentrix-2800 image file; defaults to /unix.ioe. Typically, /dev/rxdv00r:/unix (where /dev/rxdv00r is the root device).

*ddt* specifies the name of the Concentrix-2800 debugger file; defaults to /ddt.ioe. Typically, /dev/rxdv00r:/ddt (where /dev/rxdv00r is the root device).

**-c** *cfg* specifies the name of the Concentrix-2800 configuration file; defaults to /unix.cfg. Typically, /console/console.config.

**-d** *device* (option) specifies the root device, where *device* is one of the following:

*major minor*
specifies major and minor number of the root device in decimal.

*special* specifies the name of the special file for the root device, for example, /dev/rxdv00r.

Defaults to /dev/rxdv00r (major 96 and minor 17).

**-f***flags* (option) specifies boot flags: s for single-user, m for multi-user mode. If neither is specified, the console mode options are used.

**-n** (option) specifies not to load *ddt*.

**-r** *ram* (option) reserves space for a RAM disk, where *ram* is one of the following:

*sz* [*disk-image*]
specifies the amount of memory in kilobytes to reserve for the RAM disk and, optionally, the name of a file containing a disk image to load into that memory.

*header* specifies the name of a file containing a RAM disk header and, optionally, a disk image to load into memory. The header must start the file and be formatted as shown in the following example:

ramdisk file header:

```
Concentrix Ramdisk  /* must be first line in file, verbatim */
allocate=100        /* size in 'k' of memory region to allocate */
size=50             /* amount of data to read from rest of file (in k)*/
<padding>
<actual data to go in ramdisk>  /* data to go into ramdisk <if any>starts
                                   at byte 1024 in file */
```

-s (option) specifies not to start *unix* or *ddt*.

-v (option) verifies memory usage.

-V (option) prints interesting load time data.

## DESCRIPTION

This program is invoked by /console/console_loop as part of FX/2800 startup mode. You can modify the *unix*, *ddt*, and *cfg* specifications to this invocation of srboot by stopping in FX/2800 booter mode during system startup. You can specify options to this invocation of srboot by placing them in a file named /console/srbootflags under Diagnostix-2800.

## EXAMPLES

If you explicitly invoke the program, you should first initialize the system:

```
# cd /console
# dm -initialize
```

Then (staying in /console) invoke srbcot with the *unix*, *ddt*, and *cfg* specifications:

```
# srboot /dev/rxdv00r:/unix /dev/rxdv00r:/ddt -c console.config
```

To connect your console to the Concentrix-2800 system, use the tip command with sr as the argument:

```
# tip sr
```

## NOTES

This program is not intended for explicit invocation.

## SEE ALSO

makeramdisk

# srcreate

## create shared memory region

Creates a permanent shared memory region; a permanent shared region exists until explicitly deleted. Srcreate returns an error if the specified region already exists. An ordinary user cannot create a permanent region if MAX_SHARED_PERMCREATE is 0 (the default).

---

**/usr/alliant/srcreate** *id size* [*initial-value*]

---

### ARGUMENTS

*id* is an integer in the range 1 to 2**31–1 that specifies an identifier for the region.

*size* specifies the size of the region in bytes. The size must be a multiple of the system page size (4096 bytes). Permanent regions are restricted to 0 pages for ordinary users by default (the super-user is not subject to this limit); the system administrator can patch the kernel variable MAX_SHARED_PERMCREATE to change the limit. Total shared memory is restricted to the maximum of (a) 75% of available user memory or (b) available user memory minus 12MB, by default; the system administrator can change this limit with the command setmaxshare.

*initial-value*
initializes the region to the specified value; must be specified as 0 or –1.

**CAUTION:** Permanent shared memory regions use system physical memory. Very large regions can negatively impact overall system performance.

### EXAMPLE

```
# srcreate 99 4096 0
Region 99 created for 4096 bytes and initialized to 0
# srcreate 98 8192 -1
Region 98 created for 8192 bytes and initialized to -1
```

### SEE ALSO

srdelete, srdisplay
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: create_shared_region

# srdelete

## delete shared memory region

Deletes an existing, unreferenced (no attached virtual memory) shared memory region. Srdelete requires either superuser privilege or the same UID as the region's creator.

---

## /usr/alliant/srdelete *id*

---

### ARGUMENTS

*id* is an integer that identifies a shared memory region.

### EXAMPLE

```
# srdelete 99
Region 99 deleted
```

### SEE ALSO

srcreate, srdisplay
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: create_shared_region

# srdisplay

## display shared memory regions

The srdisplay command displays information about all existing shared memory regions. The information includes for each region: its identifier; the number of processes attached to the region; its size in kilobytes; the UID and GID of the creator; and the life-time of the region – permanent or temporary.

---

## /usr/alliant/srdisplay

---

### EXAMPLE

```
% srdisplay
        Id  References    Size    UID    GID    Life-Time
        99         0        8Kb      0     10    Permanent
        98         2        4Kb    116    105    Temporary
```

### SEE ALSO

srcreate, srdelete
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: create_shared_region

# statd

## network status monitor

Interacts with lockd to provide the crash and recovery functions for the locking services on NFS. Statd is an intermediate version of the status monitor.

The statd daemon is typically invoked by /etc/rc.local. (As distributed, the statd invocation is commented out in /etc/rc.local. Uncomment the invocation if you want statd to start automatically with the system.)

---

## /etc/rpc.statd

---

### NOTES

The crash of a site is only detected upon its recovery.

### FILES

/etc/sm
/etc/sm.back
/etc/state

### SEE ALSO

lockd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: statmon

# sticky

## persistent text and append-only directories

The sticky bit (file mode bit 01000 is used to indicate special treatment for certain executable files and directories. See the *Concentrix-2800 Commands and Applciations Manual,* chmod for more information.

### DESCRIPTION

While the sticky bit is set on a shareable executable file, the text of that file will not be removed from the system swap area. Thus the file does not have to be fetched from the file system upon each execution. Shareable text segments are normally placed in a least-frequently-used cache after use, so that the sticky bit has little effect on commonly-used text images.

A directory whose sticky bit is set becomes an append-only directory, or, more accurately, a directory in which the directory in which the deletion of files is restricted. A file in a sticky directory may only be removed or renamed by a user if the user has write permission for the directory and the user is the owner of the file, the owner of the directory, or the superuser. This feature is usefully applied to directories such as /tmp which must be publicly writable but should deny users the license to arbitrarily delete or rename each others' files.

Shareable executable files are made by the -n and -z options of the command ld.

### NOTES

Only the superuser can set the sticky bit on a shareable executable file.

Any user may create a sticky directory. See the chmod command for details about modifying file modes.

Since the text areas of sticky text executables are stashed in the swap area, abuse of the feature can cause a system to run out of swap.

Neither the system call open nor the system call mkdir will create a file with the sticky bit set.

### SEE ALSO

chmod, *Concentrix-2800 Commands and Applciations Manual*

# sunboot

start Sun system

Loads the kernel, or another executable program, into memory. Started from the boot PROM (console mode).

---

**>b** [*device* [(*c*,*u*,*p*)]] [*filename*] [**-a**] *boot-flags*

**>b?**

**>b!**

---

## OPTIONS

*device* is one of the following:

| | |
|---|---|
| **ie** | Intel Ethernet |
| **ec** | 3Com Ethernet |
| **le** | Lance Ethernet (Sun 3-50 system) |
| **sd** | SCSI disk |
| **st** | SCSI 1/4" tape |
| **mt** | Tape Master 9-track 1/2" tape |
| **xt** | Xylogics 1/2" tape |
| **xy** | Xylogics 440/450 disk |

*c* is the controller number; 0 if there is only one controller for the indicated type of device.

*u* is the unit number; 0 if there is only one driver.

*p* is the partition number; 0 if there is only one partition.

*filename* is the name of a stand-alone program in the selected partition, such as stand/diagi or vmunix.x. (Note: *filename* is relative to the root of the selected device and partition and never begins with a slash.) If *filename* is not specified, the boot program uses a default value (normally vmunix, which is stored in the vmunix variable in the boot executable file supplied by Sun, but can be patched with adb to indicate another stand-alone program.

**-a** prompts interactively for the device and name of the file to boot.

*boot-flags* are to the kernel or stand-alone program. They are typically arguments to that program or, as with those listed below, arguments to progams that it invokes.

**-b** passes the -b flag through the kernel to init (to skip execution of the /etc/rc.local script).

**-h** halts after loading the system.

**-s** passes the -s flag through the kernel to init (for single-user operation).

**-i** *initname* passes *initname* to the kernelas the first program to run rather than the default /sbin/init.

## DESCRIPTION

The form b? displays all boot devices and their device arguments.

The form b! boots but does not perform a RESET.

When booting stand-alone, the boot program (/boot) is brought in by the PROM from the file system. This program contains drivers for all devices.

### Booting Sun-3 Over the Network

When booting over the network, the Sun-3 system PROM obtains a version of the boot program from a server using the trivial file transfer protocol (TFTP). The client broadcasts a RARP request containing its Ethernet address. A server responds with the client's Internet address. The client then sends a TFTP request for its boot program to that server (if that fails, it broadcasts the request). The filename requested (unqualified - not a pathname) is the hexadecimal, uppercase representation of the client's Internet address, for example:

```
Using IP Address          192.9.1.17 = C0090111
```

When the Sun server receives a request, it looks in the directory /tftpboot for *filename*.e. That file is typically a symbolic link to the client's boot program, normally boot.sun3 in the same directory. The server invokes the TFTP server, tftpd, to transfer the file to the client.

When the file is successfully read by the client, the boot program jumps to the load-point and loads vmunix (or a standalone program). In order to do this, the boot program makes a broadcast RARP request to find the client's IP address, and then makes a second broadcast request to a bootparamd (bootparams daemon) for information necessary to boot the client. The bootparams daemon obtains this information either from a local /etc/bootparams database file, or from an NIS map. The boot program sends two requests to the bootparams daemon, the first, whoami, to obtain its hostname, and the second, getfile, to obtain the name of the client's server and the pathname of the client's root partition.

The boot program then performs a mount operation to mount the client's root partition, after which it can read and execute any program within that partition by pathname (including a symbolic link to another file within that same partition). Typically, the file /vmunix.x is read. If the program is not read successfully, boot responds with a short diagnostic message.

### Booting Sun-2 and Sun-4 Over the Network

Sun-2 and Sun-4 systems boot over the network in a similar fashion. However, the filename requested from a server must have a suffix that reflects the system architecture of the machine being booted. For these systems, the requested filename has the form: *ip-address.arch* where *ip-address* is the machine's Internet Protocol (IP) address in hexadecimal, and *arch* is a suffix representing its architecture. (Only Sun-3 systems can omit the arch suffix.) These filenames are restricted to 14 characters for compatibility with .UX System V and other operating systems. Therefore, the architecture suffix is limited to 5 characters; it must be in upper case. At present, the following suffixes are recognized: SUN2 for Sun-2 systems, SUN3 for Sun-3 systems, SUN4 for Sun-4 systems, and PCNFS for PC–NFS.

> NOTE: a Sun-2 system boots from its server using one extra step. It broadcasts an ND request that is intercepted by the user-level ndbootd server. This server sends back a stand-alone program that carries out the same TFTP request sequence as is done for all the other systems.

### System Start-up

Once the system is loaded and running, the kernel performs some internal housekeeping, configures its device drivers, and allocates its internal tables and buffers. The kernel then starts process number 1 to run init, which performs file system housekeeping, starts system daemons, initializes the system console, and begins multi-user operation. Some of these activities are omitted when init is invoked with certain boot flags. The boot flags are typically entered as arguments to the boot command, and passed along by the kernel to init.

## FILES

| | |
|---|---|
| /boot | standalone boot program |
| /tftpboot/???????? | symbolic link to the boot program for a client |
| /tftpboot/boot.sun3 | programs to boot from the client's root partition |
| /usr/etc/in.tftpd | TFTP server |
| /vmunix | |
| /etc/bootparams | |

## SEE ALSO

bootparamd, init, mount, ndbootd, rc, reboot, tftpd
*Concentrix-2800 Commands and Applications:* adb, tftp

# swapon

specify additional device for paging and swapping (old)

Not supported. See extendswap.

# sync

## update the super block

Writes all file system information cached in memory to disk. The information includes modified superblocks, modified i-nodes, and delayed user output.

---

## /bin/sync

---

### SEE ALSO

reboot, update
*Concentrix-2800 System Reference: System Calls and C Library Routines*: sync, fsync

# syslog

## log system messages (old)

Not supported. See syslogd.

# syslogd

## log system messages

Reads and logs messages into a set of files described by the configuration file /etc/syslog.conf. Each message is one line. A message can contain a priority code, marked by a number in angle brackets at the beginning of the line. Priorities are defined in <sys/syslog.h>. Syslogd reads from the UNIX domain socket /dev/log,from an Internet domain socket specified in /etc/services, and from the special device /dev/klog (to read kernel messages).

The syslogd daemon is typically started by the system start-up script /etc/rc.local.

---

**/etc/syslogd [-f*configfile*] [-m*markinterval*] [-d]**

---

### ARGUMENTS

**-f** specifies the configuration file; defaults to /etc/syslog.conf.

**-m** selects the number of minutes between mark messages.

**-d** turns on debugging.

### DESCRIPTION

Syslogd configures when it starts up and whenever it receives a hangup signal. Lines in the configuration file have a selector to determine the message priorities to which the line applies and an action. The action fields are separated from the selector by one or more tabs.

Selectors are semicolon-separated lists of priority specifiers. Each priority has a facility describing the part of the system that generated the message, a dot, and a level indicating the severity of the message. Symbolic names can be used. An asterisk selects all facilities. All messages of the specified level or higher (greater severity) are selected. More than one facility can be selected using commas to separate them. For example:

```
*.emerg;mail,daemon.crit
```

Selects all facilities at the emerg level, and the mail and daemon facilities at the crit level.

Known facilities and levels recognized by syslogd are those listed in the library routine syslog without the leading LOG. The additional facility mark has a message at priority LOG_INFO sent to it every 20 minutes (can be changed with the -m flag). The mark facility is not enabled by a facility field containing an asterisk. The level none can be used to disable a particular facility. For example:

```
*.debug;mail.none
```

Sends all messages except mail messages to the selected file.

The second part of each line describes where the message is to be logged if this line is selected. There are four forms:

- A filename (beginning with a leading slash). The file will be opened in append mode.
- A hostname preceded by an at sign (@). Selected messages are forwarded to syslogd on the named host.
- A comma separated list of users. Selected messages are written to those users if they are logged in.
- An asterisk. Selected messages are written to all logged-in users.

Blank lines and lines beginning with # are ignored.

For example, the configuration file:

```
kern,mark.debug             /dev/console
*.notice;mail.info          /usr/spool/adm/syslog
*.crit                      /usr/adm/critical
kern.err                    @ucbarpa
*.emerg                     *
*.alert                     eric,kridle
*.alert;auth.warning        ralph
```

Logs all kernel messages and 20-minute marks onto the system console, all notice (or higher) level messages, and all mail system messages except debug messages into the file /usr/spool/adm/syslog, and all critical messages into /usr/adm/critical; kernel messages of error severity or higher are forwarded to ucbarpa. All users are informed of any emergency messages, the users eric and kridle are informed of any alert messages, and the user ralph is informed of any alert message or any warning message (or higher) from the authorization system.

Syslogd creates the file /etc/syslog.pid, if possible, containing a single line with its process id. This file can be used to kill or reconfigure syslogd. To bring syslogd down, you should send a terminate signal (for example, kill 'cat /etc/syslog.pid').

## FILES

| | |
|---|---|
| /etc/syslog.conf | the configuration file |
| /etc/syslog.pid | the process id |
| /dev/log | name of the UNIX domain datagram log socket |
| /dev/klog | kernel log device |

## SEE ALSO

*Concentrix-2800 Commands and Applications Manual*: logger
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: syslog

# telnetd

## DARPA TELNET protocol server

Supports the DARPA standard TELNET virtual terminal protocol.

The telnet daemon is started as needed by /etc/inetd.

---

## /etc/telnetd

---

### DESCRIPTION

Telnetd allocates a pseudo-terminal device (see the special file pty) for a client, then creates a login process that has the slave side of the pseudo-terminal as stdin, stdout, and stderr. Telnetd manipulates the master side of the pseudo-terminal, implementing the TELNET protocol and passing characters between the client and the login process.

When a TELNET session is started, telnetd sends TELNET options to the client side indicating a willingness to do "remote echo" of characters, to "suppress go ahead," and to "receive terminal type information" and to "negotiate about window size." If the remote client is willing, the remote terminal type and/or window size information is propagated in the environment of the created login process. The pseudo terminal allocated to the client is configured to operate in "cooked" mode, with XTABS and CRMOD enabled (see the special file tty).

Telnetd is willing to do: echo, binary, suppress go ahead, timing mark, and negotiate about window size. Telnetd is willing to have the remote client do: binary, terminal type, suppress go ahead, and negotiate about window size.

### NOTES

- The TELNET protocol allows for the exchange of the number of lines and columns on the user's terminal, but telnetd does not make user of this feature.
- Binary mode has no common interpretation except between similar operating systems.
- The terminal type name received from the remote client is converted to lower case.
- Telnetd never sends TELNET go ahead commands.

### SEE ALSO

*Concentrix-2800 Commands and Applications*: telnet
*Concentrix-2800 System Reference: System Calls and C Library Routines*: socket
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: pty, tty, services

# tftpd

## DARPA trivial file transfer protocol server

Supports the DARPA Trivial File Transfer Protocol, operating at the port indicated in the Internet service description in the /etc/inetd.conf file.

The server is started as needed by inetd.

---

## /etc/tftpd [-s] [*homedir*]

---

### ARGUMENTS

*homedir* specifies the directory in which tftpd works. Defaults to /tftpboot.

-s (secure) requires the directory change to succeed. The daemon also changes its root directory to *homedir*.

### DESCRIPTION

The use of tftp does not require an account or password on the remote system. Due to the lack of authentication information, tftpd allows only publicly readable files to be accessed. Files may be written only if they already exist and are publicly writable. Note that this extends the concept of public to include all users on all hosts that can be reached through the network; this may not be appropriate on all systems, and its implications should be considered before enabling tftp service. The server should have the user identification with the lowest possible privilege.

Tftpd executes with the user identification and the group identification at -2, assuming that no files exist with this owner or group.

### SEE ALSO

inetd
*Concentrix-2800 Commands and Applications*: tftp
*Concentrix-2800 System Reference: System Calls and C Library Routines*: socket
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services, inetd.conf

# tnamed

## DARPA trivial name server

Supports the DARPA Name Server Protocol, operating at the port indicated in /etc/services.

The server is started as needed by inetd.

---

## /usr/etc/in.tnamed [-v]

---

### ARGUMENTS

-v (verbose) invokes the daemon in verbose mode.

### NOTES

Two known clients of this service are the MIT PC/IP software the Bridge boxes.

The protocol implemented by this program is obsolete. Its use should be phased out in favor of the Internet Domain protocol. (See named.)

### SEE ALSO

inetd
*Concentrix-2800 Commands and Applications*: uuco
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: services

# timed

## time server daemon

Synchronizes the host's time with the time of other machines in a local area network running timed.

The timed daemon is typically invoked by /etc/rc.local. (As distributed, the timed invocation is commented out in /etc/rc.local. Uncomment the invocation if you want timed to start automatically with the system.)

---

**/etc/timed [-t] [-M] [-n *network*] [-i *network*]**

---

### ARGUMENTS

**-M** permits the host to be used as a master. If the machine running the master crashes, then the slaves elect a new master from among slaves running with the -M flag. A host running without the -M flag always remains a slave.

**-t** enables timed to trace the messages it receives in the file /usr/adm/timed.log. Tracing can be turned on or off by the administration command timedc.

**-n *network*** adds the specified network to a list of valid networks; incompatible with -i. Overrides the default list of networks (see the file format networks).

**-i *network*** adds the specified network to a list of networks to be ignored; incompatible with -n. Overrides the default list of networks (see the file format networks).

### DESCRIPTION

Timed is the time server daemon and is normally invoked at boot time from the rc file. Time servers slow down the clocks of some machines and speed up the clocks of others to bring them to the average network time. The average network time is computed from measurements of clock differences using the ICMP timestamp request message.

The service provided by timed is based on a master-slave scheme. When timed is started on a machine, it asks the master for the network time and sets the host's clock to that time. After that, it accepts synchronization messages periodically sent by the master and calls the system call adjtime to perform the needed corrections on the host's clock.

Timed also communicates with the command date in order to set the date globally, and with the administration command timedc, a timed control program.

Timed normally checks for a master time server on each network to which it is connected, except as modified by the options -n and -i. It requests synchronization service from the first master server located. If permitted by the -M flag, it provides synchronization service on any attached networks on which no current master server is detected. Such a server propagates the time computed by the top-level master.

## FILES

| | |
|---|---|
| /usr/adm/timed.log | tracing file for timed |
| /usr/adm/timed.masterlog | log file for master timed |

## SEE ALSO

timedc
*Concentrix-2800 Commands and Applications*: date
*Concentrix-2800 System Reference: System Calls and C Library Routines*: adjtime, gettimeofday
*Concentrix-2800 System Reference: Miscellaneous* Routines, *Special Files, and File Formats*: icmp

# timedc

## timed control program

Controls the operation of the timed program.

---

**/etc/timedc** [*command* [*argument* ...]]

---

### DESCRIPTION

Timedc can be used to:

- measure the differences between machines' clocks.
- find the location where the master time server is running.
- enable or disable tracing of messages received by timed.
- perform various debugging actions.

Without arguments, timedc prompts for commands from the standard input. If arguments are supplied, timedc interprets the first argument as a command and the remaining arguments as parameters to the command. The standard input can be redirected causing timedc to read commands from a file. Commands can be abbreviated.

**?** [*command* ...]
: prints a short description of each command specified in the argument list, or, if no arguments are specified, a list of the recognized commands.

**help** [*command* ...]
: prints a short description of each command specified in the argument list, or, if no arguments are specified, a list of the recognized commands.

**clockdiff** *host* ...
: computes the differences between the clock of the host machine and the clocks of the specified machines.

**trace** {**on**|**off**}
: enables or disables the tracing of incoming messages to timed in the file /usr/adm/timed.log.

**quit**
: exits from timedc.

Other commands can be included for use in testing and debugging timed; the help command can be consulted for details.

## FILES

| | |
|---|---|
| /usr/adm/timed.log | tracing file for timed |
| /usr/adm/timed.masterlog | log file for master timed |

## SEE ALSO

timed
*Concentrix-2800 Commands and Applications*: date
*Concentrix-2800 System Reference: System Calls and C Library Routines*: adjtime
*Concentrix-2800 System Reference: Miscellaneous* Routines, *Special Files, and File Formats*: icmp

## DIAGNOSTICS

?Ambiguous command
: abbreviation matches more than one command

?Invalid command
: no match found

?Privileged command
: command can be executed by root only

# timezone

## set time zone

The timezone command sets the time zone as specified. As shipped, the date in Concentrix-2800 is set to Eastern US time (EST or EDT).

**NOTE:** Execute this command in single-user mode only.

---

**/usr/alliant/timezone** *option ...*

---

### ARGUMENTS

| | |
|---|---|
| **et** | Eastern US Time |
| **ct** | Central US Time |
| **mt** | Mountain US Time |
| **pt** | Pacific US Time |
| **mine** *n* | minutes east of Greenwich |
| **minw** *n* | minutes west of Greenwich |
| **nodst** | standard time (not on Daylight Saving Time) |
| **us** | US Daylight Saving Time |
| **aus** | Australian Daylight Saving Time |
| **we** | Western European Daylight Saving Time |
| **me** | Mid European Daylight Saving Time |
| **ee** | East European Daylight Saving Time |

### EXAMPLE

For example, to set the time zone for California:

```
# /usr/alliant/timezone pt us
```

For example, to set the time zone for Paris:

```
# /usr/alliant/timezone mine 60 me
```

### NOTES

- Timezone is a shell script that invokes adb to patch the Concentrix-2800 kernel, both on disk and in memory.

### SEE ALSO

*Concentrix-2800 Commands and Applications*: date

# trpt

## transliterate protocol trace

This software is provided on an "as is" basis without warranty expressed or implied. Any software services available will be provided at the then current charges.

Interrogates the buffer of TCP trace records created when a socket is marked for "debugging" (see setsockopt), and prints a readable description of these records.

---

**/etc/trpt** [*option ...*] [*system* [*core*]]

---

### ARGUMENTS

By default, trpt prints all the trace records found in the system grouped according to TCP connection protocol control block (PCB). The following options can be used to alter this behavior:

**-a** prints the values of the source and destination addresses for each packet recorded, in addition to the normal output.

**-f** follows the trace as it occurs, waiting a short time for additional records each time the end of the log occurs.

**-j** gives just a list of the PCB addresses for which there are trace records.

**-p***address* shows only trace records associated with the PCB (hexadecimal *address* follows).

**-s** prints a detailed description of the packet sequencing information, in addition to the normal output.

**-t** prints the values for all timers at each point in the trace, in addition to the normal output.

*system* specifies a non-default system file for debugging.

*core* specifies a non-default core file for debugging.

### DESCRIPTION

The recommended use of trpt is as follows. Isolate the problem and enable debugging on the socket(s) involved in the connection. Find the address of the PCBs associated with the sockets using the -A option to netstat. Then run trpt with the -p option, supplying the associated PCB addresses. The -f option can be used to follow the trace log once the trace is located. If there are many sockets using the debugging option, the -j option may be useful in checking to see if any trace records are present for the socket in question.

## FILES

/vmunix
/dev/kmem

## DIAGNOSTICS

- `no namelist` – The system image doesn't contain the proper symbols to find the trace buffer; others messages should be self explanatory.

## NOTES

- Trpt should also print the data for each input or output, but this is not saved in the race record.
- The output format is inscrutable and should be described here.

## SEE ALSO

*Concentrix-2800 Commands and Applications*: netstat
*Concentrix-2800 System Reference: System Calls and C Library Routines*: setsockopt

# trsp

## transliterate sequenced packet protocol trace

Interrogates the buffer of SPP trace records created when a socket is marked for debugging (see the system call setsockopt), and prints a readable description of these records.

---

**/etc/trsp [-a] [-s] [-t] [-j] [-p *hex-address*] [*system* [*core*]]**

---

### ARGUMENTS

When no options are supplied, trsp prints all the trace records found in the system grouped according to SPP connection protocol control block (PCB).

**-a** in addition to the normal output, prints the values of the source and destination addresses for each packet recorded.

**-j** just gives a list of the protocol control block addresses for which there are trace records.

**-p** shows only trace records associated with the protocol control block whose address follows.

**-s** in addition to the normal output, prints a detailed description of the packet sequencing information.

**-t** in addition to the normal output, prints the values for all timers at each point in the trace.

### DESCRIPTION

The recommended use of trsp is as follows. Isolate the problem and enable debugging on thesocket(s) involved in the connection. Find the address of the protocol control blocks associated with the sockets using the -A option to netstat.Then run trsp with the -p option, supplying the associated protocol control block addresses. If there are many sockets using the debugging option, the -j option may be useful in checking to see if any trace records are present for the socket in question.

If debugging is being performed on a system or core file other than the default, the last two arguments may be used to supplant the defaults.

### DIAGNOSTICS

no namelist
: the system image does not contain the proper symbols to find the trace buffer.

Others are self-explanatory.

## FILES

| | |
|---|---|
| /vmunix | UNIX kernel image |
| /dev/kmem | kernel memory |

## SEE ALSO

*Concentrix-2800 Commands and Applications*: netstat
*Concentrix-2800 System Reference: System Calls and C Library Routines*: setsockopt

# tunefs

## tune up an existing file system

Changes some of the dynamic parameters of a file system that affect the layout policies.

---

**/etc/tunefs** *option ... filesys*

---

### ARGUMENTS

**-a** *n* specifies the maximum number of contiguous blocks laid out before forcing a rotational delay (see -d below). The default value is 1024; Concentrix-2800 disk drivers are capable of reading or writing many contiguous blocks in a single transfer.

This option may not be specified on TFS filesystems.

**-A** copies the superblock to all alternate superblocks after a dynamic parameter change. Another option must also be specified.

**-d** *n* specifies the expected time (in milliseconds) to service a transfer completion interrupt and initiate a new transfer on the same disk. It is used to decide how much rotational spacing to place between successive blocks in a file. The default is 5 for all supported devices. Because of the -a default value of 1024, this only matters when blocks cannot be allocated sequentially.

This option may not be specified on TFS filesystems.

**-e** *n* indicates the maximum number of blocks any single file can allocate out of a cylinder group before it is forced to begin allocating blocks from another cylinder group. Typically this value is set to about one quarter of the total blocks in a cylinder group. The intent is to prevent any single file from using up all the blocks in a single cylinder group, thus degrading access times for all files subsequently allocated in that cylinder group. The effect of this limit is to cause big files to do long seeks more frequently than if they were allowed to allocate all the blocks in a cylinder group before seeking elsewhere. For file systems with exclusively large files, this parameter should be set higher.

This option may not be specified on TFS filesystems.

**-m** *n* specifies the percentage of space held back from normal users; the minimum free space threshold. The default value used is 10%. This value can be set to zero; however up to a factor of three in throughput is lost over the performance obtained at a 10% threshold. Note that if the value is raised above the current usage level, users will be unable to allocate files until enough files have been deleted to get under the higher threshold.

**-o** *space* or *time*
specifies whether fragmnet allocation is to be optimized for *space* or *time*. The file system adjusts this value based on the degree of disk fragmentation. Manual adjustment is not required.

This option may not be specified on TFS filesystems.

*filesys* name of the file system. The name can be the mount point, the raw device name, or the block device name.

## NOTES

- Unmount a file system before tuning it.
- If /etc/tunefs is executed on the root file system, the system must be rebooted.
- Changing -a or -d options on a disk with significant space allocated may not have the desired effect, since the pattern of free space will reflect the old parameters for a considerable time.

## SEE ALSO

dumpfs, makefs
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: fs

# umount

## unmount file systems

Unmounts a file system from a directory (the mount point).

---

**/etc/umount [-v] [-t *type*] [-h *host*]**

**/etc/umount -a[v]**

**/etc/umount [-v] *fsname-or-dir* ...**

---

### ARGUMENTS

**-a** unmounts all file systems currently mounted as determined from /etc/mtab.

**-h** *host*
unmounts all file systems currently mounted from the specified host as determined from /etc/mtab.

**-t** *type*
unmounts all file systems currently mounted as the specified type as determined from /etc/mtab.

**-v** (verbose) displays a message indicating the filesystem being unmounted.

*fsname* specifies the name of a file system.

*dir* specifies the name of the mount point.

Specify the file system name or the mount point, not both.

### DIAGNOSTICS

- `Mount device busy` – The file system is currently in use or a user has the current working directory set in that file system.

### EXAMPLES

Unmounts all file systems:

```
# /etc/umount -a
```

Unmounts the file system at the mount point /px:

```
# /etc/umount /px
```

### FILES

| | |
|---|---|
| /etc/mtab | mounted file system table |
| /etc/fstab | static information about file systems |

## SEE ALSO

mount, mountd, nfsd
*Concentrix-2800 System Reference: System Calls and C Library Routines*: mount, nfsmount, umount
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: fstab, mtab

# unmkstripe

## disassemble disk stripe

Releases disk partitions from participation in a disk stripe.

---

**/etc/unmkstripe -a**

**/etc/unmkstripe** *raw-stripe*

---

*raw-stripe* is the raw device name of the disk stripe.

**-a** unmakes all disk stripes described in /etc/stripetab.

### EXAMPLES

The following example unmakes a disk strip:

```
# /etc/unmkstripe /dev/rstripe0
```

The following example unmakes all disk stripes:

```
# /etc/unmkstripe -a
```

### FILES

/etc/stripetab disk stripe data base

### SEE ALSO

mkstripe
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: stripetab

# update

## periodically update the superblock

Performs the same function as /etc/sync every 30 seconds, keeping the file system fairly up to date in case of a crash.

The update command is typically invoked by rc.

---

## /etc/update

---

### SEE ALSO

sync, init, rc
*Concentrix-2800 System Reference: System Calls and C Library Routines*: sync

# uucico

## transfer files queued by uucp or uux

Transfers files between systems. Uucp and uux queue requests for the data transfers.

---

**/usr/lib/uucp/uucico** [**-d***spooldir*] [**-g***grade*] [**-r***role*] [**-R**] [**-s***system*] [**-x***debug*] [**-L**] [**-t***turnaround*]

---

### ARGUMENTS

**-d***spooldir*
uses spooldir as the spool directory. The default is /usr/spool/uucp.

**-g***grade*
only sends jobs of grade grade or higher this transfer. The grade of a job is specified when the job is queued by uucp or uux.

**-r***role* is either 1 or 0 and indicates whether uucico is to start up in master or slave role, respectively. A value of 1 is used when running uucico by hand or from cron. A value of 0 is used when another system calls the local system. Slave role is the default.

**-R** reverses roles. When used with the -r1 option, tells the remote system to begin sending its jobs first, instead of waiting for the local machine to finish.

**-s***system*
calls only system system. If -s is not specified and -r1 is specified, uucico attempts to call all systems for which there is work. If -s is specified, a call is made even if there is no work for that system. Useful for polling.

**-x***debug*
turns on debugging at level *debug*. Level 5 is a good start when trying to find out why a call failed. Level 9 is very detailed. Level 99 is very verbose. If *role* is 1 (master), output is normally written to the standard error output stderr. If stderr is unavailable, output is written to /usr/spool/uucp/AUDIT/system. When role is 0 (slave), debugging output is always written to the AUDIT file.

**-L** only calls local sites. A site is considered local if the device-type field in L.sys is one of LOCAL, DIR, or TCP.

**-t***turnaround*
use *turnaround* as the line turnaround time (in minutes) instead of the default 30. If turnaround is missing or 0, line turnaround is disabled. After uucico has been running in slave role for *turnaround* minutes, it attempts to run in master role by negotiating with the remote machine. In earlier versions of uucico, a transfer of many large files in one direction would hold up mail going in the other direction. With the turnaround code working, the message flow is more bidirectional in the short term.

A SIGFPE signal (see kill) to uucico toggles the debugging on or off.

Uucico is commonly used either of two ways: as a daemon run periodically by cron to call out to remote systems, and as a shell for remote systems who call in. For calling out periodically, a typical line in crontab would be as follows:

```
0   *    *  *    *    /usr/lib/uucp/uucico -r1
```

This line runs uucico every hour in master role. For each system that has transfer requests queued, uucico calls the system, logs in, and executes the transfers. The file L.sys is consulted for information about how to log in, while L-devices specifies available lines and modems for calling.

For remote systems to dial in, an entry in the passwd file must be created, with a login shell of uucico. For example:

```
ucp:Password:6:1::/usr/spool/uucppublic:/usr/lib/uucp/uucico
```

The UID for UUCP remote logins is not critical, so long as it differs from the UUCP administrative login. The latter owns the UUCP files, and assigning this UID to a remote login would be an extreme security hazard.

**FILES**

| | |
|---|---|
| /usr/lib/uucp/ | UUCP internal files/utilities |
| /usr/lib/uucp/L-devices | Local device descriptions |
| /usr/lib/uucp/L-dialcodes | Phone numbers and prefixes |
| /usr/lib/uucp/L.aliases | Hostname aliases |
| /usr/lib/uucp/L.cmds | Remote command permissions list |
| /usr/lib/uucp/L.sys | Host connection specifications |
| /usr/lib/uucp/USERFILE | Remote directory tree permissions list |
| | |
| /usr/spool/uucp/ | Spool directory |
| /usr/spool/uucp/AUDIT/* | Debugging audit trails |
| /usr/spool/uucp/C./ | Control files directory |
| /usr/spool/uucp/D./ | Incoming data file directory |
| /usr/spool/uucp/D.hostname/ | Outgoing data file directory |
| /usr/spool/uucp/D.hostnameX/ | Outgoing execution file directory |
| /usr/spool/uucp/CORRUPT/ | Place for corrupted C. and D. files |
| /usr/spool/uucp/ERRLOG | UUCP internal error log |
| /usr/spool/uucp/LOGFILE | UUCP system activity log |
| /usr/spool/uucp/LCK/LCK..* | Device lock files |
| /usr/spool/uucp/SYSLOG | File transfer statistics log |
| /usr/spool/uucp/STST/* | System status files |
| /usr/spool/uucp/TM./ | File transfer temp directory |
| /usr/spool/uucp/X./ | Incoming execution file directory |
| | |
| /usr/spool/uucppublic | Public access directory |

## SEE ALSO

uuclean, uucpd, uupoll, uusnap, uuxqt
*Concentrix-2800 Commands and Applications* – uucp, uuq, uux
*Concentrix-2800 System Reference, Miscellaneous Routines, Special Files, and File Formats:* L-devices, L-dialcodes, L.aliases, L.cmds, L.sys

# uuclean

## UUCP spool directory clean-up

**Uuclean** scans the spool directory for files with the specified prefix and deletes all those that are older than the specified number of hours.

---

## /usr/lib/uucp/uuclean [*option*]

---

### ARGUMENTS

**–p***pre* scans for files with *pre* as the file prefix. Up to ten –p arguments can be specified. A –p without any *pre* following causes all files older than the specified time to be deleted.

**–n***time* deletes files whose age is more than *time* hours if the prefix test is satisfied. (default time is 72 hours)

**–m** sends mail to the owner of the file when it is deleted.

This program is typically started by *cron*.

### FILES

| | |
|---|---|
| /usr/lib/uucp | directory with commands used by uuclean internally |
| /usr/lib/uucp/spool | spool directory |

### SEE ALSO

*Concentrix-2800 Commands and Applications*: uucp, uux

# uucpd

## UUCP daemon

Supports uucp connections over networks.

---

## /etc/uucpd

---

### DESCRIPTION

Uucpd listens for service requests at the port indicated in the service name data base; see the services special file. The server provides login name and password authentication before starting up uucico for the rest of the transaction.

| | |
|---|---|
| /usr/spool/uucppublic | Public access directory |

### SEE ALSO

uucico, uuclean, uupoll, uusnap, uuxqt
*Concentrix-2800 Commands and Applications*: uucp, uuq, uux
*Concentrix-2800 System Reference, Miscellaneous Routines, Special Files, and File Formats*: L-devices, L-dialcodes, L.aliases, L.cmds, L.sys, services

# uupoll

## poll a remote UUCP site

Forces a poll of a remote system. Uupoll queues a null job for the remote system and invokes uucico.

---

**/usr/bin/uupoll [-g*grade*] [-n] *system***

---

-g*grade*
: only sends jobs of grade *grade* or higher on this call.

-n
: queues the null job, but does not invoke uucico.

### DESCRIPTION

Uupoll is usually run by cron or by a user who wants to hurry a job along. A typical entry in crontab:

```
0    0,8,16    *    *    *    /usr/bin/uupoll ihnp4
0    4,12,20   *    *    *    /usr/bin/uupoll ucbvax
```

This entry polls ihnp4 at midnight, 0800, and 1600, and ucbvax at 0400, noon, and 2000.

If the local machine is already running uucico every hour and has a limited number of outgoing modems, a more elegant approach might be:

```
0    0,8,16    *    *    *    /usr/bin/uupoll -n ihnp4
0    4,12,20   *    *    *    /usr/bin/uupoll -n ucbvax
5    *         *    *    *    /usr/lib/uucp/uucico -rl
```

This entry queues null jobs for the remote sites at the top of the hour; the jobs are processed by uucico when it runs five minutes later.

### FILES

/usr/lib/uucp/ UUCP internal files/utilities
/usr/spool/uucp/ Spool directory

### SEE ALSO

uucp, uux, uucico

# uusnap

## show snapshot of the UUCP system

Uusnap displays in tabular format a synopsis of the current UUCP situation.

---

## /usr/bin/uusnap

---

### DESCRIPTION

The format of each line is as follows:

site N Cmds N Data N Xqts Message

*site* is the name of the site with work,

*N* is a count of each of the three possible types of work (command, data, or remote execute), and

*Message* is the current status message for that site as found in the STST file.

Included in *Message* may be the time left before UUCP can re-try the call, and the count of the number of times that UUCP has tried to reach the site.

### SEE ALSO

*Concentrix-2800 Commands and Applications*: uucp

# uuxqt

## UUCP execution file interpreter

Interprets execution files created on a remote system via uux and transferred to the local system via uucico.

---

**/usr/lib/uucp/uuxqt** [**-x***debug*]

---

### DESCRIPTION

When a user uses uux to request a remote command, uuxqt actually executes the command. Normally, uuxqt is forked from uucico to process queued execution files; for debugging, it may also be run manually by the UUCP administrator.

Uuxqt runs in its own subdirectory /usr/spool/uucp/XTMP and copies intermediate files to this directory when necessary.

### FILES

| | |
|---|---|
| /usr/lib/uucp/L.cmds | Remote command permissions list |
| /usr/lib/uucp/USERFILE | Remote directory tree permissions list |
| /usr/spool/uucp/LOGFILE | UUCP system activity log |
| /usr/spool/uucp/LCK/LCK.XQT | Uuxqt lock file |
| /usr/spool/uucp/X./ | Incoming execution file directory |
| /usr/spool/uucp/XTMP | Uuxqt running directory |

### SEE ALSO

uucico, uucp, uux,
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: L.cmds, USERFILE

# uv

## print kernel version information

Prints to standard output the version name, version date, author, version number (in hexadecimal), and version patch (in hexadecimal) of a Concentrix-2800 kernel image.

---

**/usr/alliant/uv** [*kernel-name*]

---

### ARGUMENTS

*kernel-name*
(defaults to /unix) name of the kernel image.

### EXAMPLE

```
# /usr/alliant/uv
._versionname:  Alliant Concentrix 2800 Rel. 1.0.06
._versiondate:  built Sun May 6 2:56:22 EDT 1990
._versionmake:  by root (/src1/Sunrise/sr1src/usr/src/sys/CONCENTRIX)
._version:      version# 0x10060002
._versionpatch: versionpatch# 0x0
```

# verify_tapevol_arch

verify mounted distribution tape is correct

Verifies that the software distribution tape loaded is the proper one. It is called by setup_exec as part of the loading of diskless/NFS architecture support to verify that the requested tape is indeed loaded. If the mounted tape is not correct, the script prompts the user to load the correct tape.

---

**/usr/etc/install/verify_tapevol_arch** *arch volume tapedev* [*tapeserver*]

---

## ARGUMENTS

*arch* is the architecture name of the desired distribution tape, typically one of the values sun2, sun3, or sun4.

*volume* is the desired volume (tape) number of the distribution.

*tapedev* is the full pathname of the non-rewinding raw tape device on which the tape is loaded. This device may be on a remote machine. If the device is on a remote machine, *tapeserver* must be specified as the host name of the remote machine with the tape device.

## DESCRIPTION

The script first rewinds the loaded tape. The second tape file (the tape's table of contents in XDR format) is read and piped through xdrtoc. This output is left in the file /tmp/TOC. If the tape's volume number and architecture type are not correct, a message is output to the user to mount the correct tape.

## FILES

/tmp/TOC

## SEE ALSO

INSTALL, extracting, setup_exec, xdrtoc

# vipw

## edit the password file

Locks the file /etc/passwd and invokes an editor on the file. After you exit from the editor, builds the data base files /etc/passwd.pag and /etc/passwd.dir, and unlocks /etc/passwd.

---

## /etc/vipw

---

### DESCRIPTION

If the password file is already being edited, you are told to try again later. The vi editor is used unless the environment variable EDITOR indicates an alternate editor. Vipw performs a number of consistency checks on the password entry for root to prevent installation of a password file with a mangled root entry.

### FILES

| | |
|---|---|
| /etc/passwd | password file |
| /etc/passwd.dir | password file in dbm format |
| /etc/passwd.pag | password file in dbm format |
| /etc/ptmp | temporary file |

### SEE ALSO

*Concentrix-2800 Commands and Applications*: chsh, mkpasswd, passwd
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: passwd

# vmlock

## enable/disable process swapping and outpaging

Controls swapping and working set management for a process. Superuser privilege is required.

---

**/usr/alliant/vmlock -e** *pid*

**/usr/alliant/vmlock -d** *pid*

---

### ARGUMENTS

**-e** enables virtual memory locking, that is, locks the specified process into memory by inhibiting swapping and paging.

**-d** disables virtual memory locking, that is, permits process swapping and paging.

*pid* the process id of the process to be affected.

### SEE ALSO

*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: enable_vmlock

# xdrtoc

## convert distribution table of contents

Translates the XDR-format table of contents found on software distribution tapes into a human-readable form on standard output.

---

**/usr/etc/install/xdrtoc** [*file*]

---

### ARGUMENTS

*file* specifies the XDR table of contents; defaults to standard input.

### FILES

/tmp/TOC human readable form of the XDR table of contents

### SEE ALSO

INSTALL, setup_exec, verify_tapevol_arch

# xdvstat

## xdv disk statistics dynamic display

The xdvstat command gives a continuous display of statistics maintained by the xdv disk driver. These statistics are kept if the system is booted with the kernel longword xdv_statistics set at non-zero.

---

**/usr/alliant/xdvstat [-i#] [-b]**

---

### ARGUMENTS

**-i#** pause for interval of # seconds between displays of the current disk statistics. Default is 10 seconds.

**-b** brief display. Limit the display to a 24-line screen type display.

### DESCRIPTION

xdvstat first checks the size of the screen of the output device, then determines which xdv disks are online, and displays as many as will fit on the screen. xdvstat displays a "brief" screen unless there are at least 34 lines on the terminal. See the Statistics section below for details about which items are displayed.

Use the **-i#** command to set the interval between statistics retrieval. Bear in mind that these statistics are maintained by the xdv driver code within the kernel; xdvstat merely collects them at each interval. You can alter the disks displayed by deleting disks with the **d** command, then adding some with the **a** command. You can use **r** to zero certain statistic counts, to make the display easier to read.

xdvstat acts upon single keys typed at the control terminal to allow the user to change its display and/or operating parameters. Note that some commands request additional input (at least a RETURN to continue the display). You can change the arguments of xdvstat while it is running.

Type CTRL C or **q** at any point to quit xdvstat.

### Command Keys

**a, A, or +** Add a disk to the display (enter two digits). For example, to add xdv43 to the display, enter:

a 43 RETURN

**b or B** Brief display. Like the **-b** command line option; treat the display as if there were only 24 lines on the screen.

**d, D, or -** Delete disk from display (enter two digits). For example, to delete xdv12 from the display, enter:

d 12 RETURN

| | |
|---|---|
| **f or F** | Full display. Show additional information, if it fits. (If >33 rows) |
| **h, H, or ?** | Display Help screen. Requires RETURN to continue. |
| CTRL L | Redraw the screen. |
| **i or I** | Set sample Interval. Enter number of seconds and RETURN. |
| *** or .** | Display all disks that fit (depending on screen size). |
| **o or O** | List Online disks. |
| **r or R** | Reset counts to zero. Sets certain statistical values to zero, so that they're more easily tracked as they grow. |
| **t or T** | Display Totals since boot. This command undoes the Reset command. |
| **q or Q** | quit xdvstat. |

**Controller Statistics**

Items marked with a * are reset to zero by the r command. Items marked with a + are included in the (b)rief display:

```
idle_cnt     * +   Number of times controller was idle
act_cnt      * +   Number of times controller was active
timestamp          Start of busy controller time
act_time     * +   Total active time
iopb_gfail   * +   Iopb allocation failure (global)
iopb_lfail         Iopb allocation failure (local)
iopb_maxq      +   Maximum number of outstanding iopbs
deq_time           Last iopb deque time
```

**Drive Statistics**

Items marked with a * are reset to zero by the r command. Items marked with a + are included in the (b)rief display:

```
write_cnt    * +   Number of successful write operations
minwrblks      +   Minimum blocks written in one operation
maxwrblks      +   Maximum blocks written in one operation
minwrqtime         Minimum time in queue for write
maxwrqtime         Maximum time in queue for write
totwrqtime         Total time in queue for write
minwratime         Minimum active time for write
maxwratime         Maximum active time for write
totwratime         Total active time for write
read_cnt     * +   Number of successful read operations
minrdblks      +   Minimum blocks read in one operation
maxrdblks      +   Maximum blocks read in one operation
minrdqtime         Minimum time in queue for read
maxrdqtime         Maximum time in queue for read
totrdqtime         Total time in queue for read
minrdatime         Minimum active time for read
```

```
maxrdatime          Maximum active time for read
totrdatime          Total active time for read
misc_cnt    * +     Miscellaneous commands
err_cnt     * +     Number of unsuccessful operations
retry_cnt   * +     Number of retries
```

## EXAMPLE

The following example shows a brief display.

```
columbo%
xdvstat -b columbo                                        Wed Jun 19 17:01:55 1991

Controller->      xdv0*                 xdv1*                 xdv2*
idle_cnt*...      12724                  4577                  2158
act_cnt*....       2622                     0                  4255
act_time*...0x01c57e0f           0x00209b15            0x00899f18
iopb_gfail*.        705                     0                     0
iopb_maxq... 20 of  20           2 of  20              4 of  20
Disk Unit-->      xdv00      xdv01      xdv10      xdv11      xdv20      xdv21
write_cnt*..       1080         19          2          0         11          2
minwrblks...          1          1          4         -1          4          2
maxwrblks...          8          8          8          0          8          4
totwrblks*..       8474        125         12          0         84          6
read_cnt*...       6206       8041       2299       2276       2304       4096
minrdblks...          2          2          2          2          2          2
maxrdblks...          8          8         88          8         64          8
totrdblks*..      43000      50770      18262      17994      18310      31062
misc_cnt*...          0          0          0          0          0          0
error_cnt*..          0          0          0          0          0          0
retry_cnt*..          0          0          0          0          0          0
          Type '?' for help
```

# ypbind

## network information service binder process

Associates the current domain with the internet address of the NIS server and the port on the host at which the /etc/ypserv process is listening for service requests. The ypbind daemon must be running on every host using NIS pages.

The ypbind daemon is typically invoked by /etc/rc.local. You must edit the line that invokes /bin/domainname and replace mydomain with the name of your domain.

---

## /usr/etc/ypbind [-s] [-ypset|-ypsetme]

---

### ARGUMENTS

| | |
|---|---|
| **-s** | Secure. When specified, only ypservers bound to a reserved port are used. This allows for a slight increase in security in completely controlled environments, where there are no computers operated by untrusted individuals. It offers no real increase in security. |
| **-ypset** | The ypset command may be used to change the binding. This option is very dangerous, and only should be used for debugging the network from a remote machine. |
| **-ypsetme** | ypset may be issued from this machine, security is based on IP address checking, which can be defeated on network where untrusted individuals may inject packets. This option is not recommended. |

### DESCRIPTION

The function of ypbind is to remember information that lets client processes on a single node communicate with some ypserv process. Ypbind must run on every machine that has NIS client processes; ypserv may or may not be running on the same node, but must be running somewhere on the network.

The information ypbind remembers is called a binding – the association of a domain name with the internet address of the NIS server, and the port on that host at which the ypserv process is listening for service requests. This information is cached in the directory /etc/yp/binding using a filename of domainname.version.

The process of binding is driven by client requests. As a request for an unbound domain comes in, the ypbind process broadcasts on the net trying to find a ypserv process that serves maps within that domain. Since the binding is established by broadcasting, there must be at least one ypserv process on every net. Once a domain is bound by a particular ypbind, that same binding is given to every client process on the node. The ypbind process on the local node or a remote node may be queried for the binding of a particular domain by using the ypwhich command.

Bindings are verified before they are given out to a client process. If ypbind is unable to speak to the ypserv process it is bound to, it marks the domain as unbound, tells the client process that the domain is unbound, and tries to bind the domain once again. Requests received for an unbound domain fail immediately. In general, a bound domain is marked as unbound when the node running ypserv crashes or gets overloaded. In such a case, ypbind binds to any NIS server (typically one that is less heavily loaded) available on the net.

Ypbind also accepts requests to set its binding for a particular domain. The request is usually generated by the NIS subsystem itself. Ypset is a command to access the Set_domain facility. It is for unsnarling messes, not for casual use. Note that the Set Domain procedure only accepts requests from processes running as root.

## FILES

The file(s) /etc/yp/binding/domainname.version will be created to speed up the binding process. These files cache the last successful binding created for the given domain; when a binding is requested these files are checked for validity and then used.

/etc/yp
/usr/etc/ypbind

## SEE ALSO

makedbm, ypinit, ypmake, yppoll, yppush, ypserv, ypset, ypwhich, ypxfr
*Concentrix-2800 Commands and Applications:* domainname, ypcat, ypmatch
*Concentrix-2800 System Reference*: *System Calls and C Library Routines:* dbm, ypclnt
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# ypinit

## build and install network information service data base

Sets up a network information service (NIS) data base on a server machine. The server can be a master server or a slave server. Requires superuser privilege.

---

**/usr/etc/yp/ypinit -m**

**/usr/etc/yp/ypinit -s** *master-name*

---

### ARGUMENTS

**-m** indicates that the local host is to be the NIS master.

**-s** sets up a slave database.

### DESCRIPTION

Ypinit asks a few self-explatory questions and reports success or failure to the terminal.

Ypinit sets up a **master** server using the simple model in which that server is master to all maps in the data base. Later you can change the association of maps to masters by changing the information in the NIS map (by applying /etc/yp/makedbm to ypmaps). All databases are built from scratch, either from information available to the program at runtime, or from the ASCII data base files in /etc. The files used to generate the data bases are listed below under FILES; see ypmake for a more complete description. The base files should not be in the abbreviated form used on client machines.

Ypinit sets up a **slave** server by copying an existing data base from a running server. The *master-name* argument specifies the host name of NIS server. If all maps have one master and that server is up, specify the host name of that server. Otherwise, specify the host name of any stable server that is up.

### FILES

/etc/passwd
/etc/group
/etc/hosts
/etc/networks
/etc/services
/etc/protocols
/etc/ethers

### SEE ALSO

makedbm, ypmake, yppush, ypserv, ypxfr
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# ypmake

## rebuild network information sevice data base

Executes the makefile for building an NIS database for a domain. The makefile generates NIS maps from base files on the host system and optionally pushes those maps to other servers in the domain. The makefile also creates empty files to time stamp the operations.

---

**cd /etc/yp ; make** [*map*]

---

### VARIABLES

The makefile contains three special variables. Edit /etc/yp/Makefile to change the defaults.

DIR — specifies the directory containing the base files. Defaults to /etc.

DOM — specifies the domain name. The domain must exist as a subdirectory of the main NIS directory (by default, /etc/yp). Defaults to the current domain.

NOPUSH — if NOPUSH contains any value, does not push new database files to other servers in the domain. Defaults to a null string.

ALIASES — specifies the aliases base file. Defaults to /usr/lib/aliases.

### ARGUMENTS

With no arguments, the makefile creates dbm databases for all maps that are out of date. An argument makes only the specified map (if it is out of date):

passwd — generates the following files from $(DIR)/passwd:

```
etc/yp/$(DOM)/passwd.byname
etc/yp/$(DOM)/passwd.byuid
etc/yp/passwd.time
```

group — generates the following files from $(DIR)/group:

```
etc/yp/$(DOM)/group.byname
etc/yp/$(DOM)/group.byuid
etc/yp/group.time
```

hosts — generates the following files from $(DIR)/hosts:

```
etc/yp/$(DOM)/hosts.byname
etc/yp/$(DOM)/hosts.byaddr
etc/yp/hosts.time
```

ethers generates the following files from $(DIR)/ethers:

etc/yp/$(DOM)/ethers.byname
etc/yp/$(DOM)/ethers.byaddr
etc/yp/ethers.time

networks generates the following files from $(DIR)/networks:

etc/yp/$(DOM)/networks.byname
etc/yp/$(DOM)/networks.byaddr
etc/yp/networks.time

services generates the following files from $(DIR)/services:

etc/yp/$(DOM)/services.byname
etc/yp/services.time

rpc generates the following files from $(DIR)/rpc:

etc/yp/$(DOM)/rpc.byname
etc/yp/rpc.time

protocols generates the following files from $(DIR)/protocols:

etc/yp/$(DOM)/protocols.byname
etc/yp/$(DOM)/protocols.bynumber
etc/yp/protocols.time

netgroup generates the following files from $(DIR)/netgroup:

etc/yp/$(DOM)/netgroup
et/yp/$(DOM)/byuser
et/yp/$(DOM)/byhost
etc/yp/services.time

aliases generates the following files from $(ALIASES):

etc/yp/$(DOM)/mail.aliases
etc/yp/aliases.time

usrclass generates the following files from $(DIR)/usrclass:

etc/yp/$(DOM)/usrclass.byname
etc/yp/$(DOM)/usrclass.byuid
etc/yp/usrclass.time

## SEE ALSO

makedbm, yppush, ypserv
*Concentrix-2800 Commands and Applications:* domainname, make

# yppasswdd

server for modifying network information service password file

Handles password change requests from the general command yppasswd.

---

**/usr/etc/rpc.yppasswdd** *file* [*adjunct_file*] [**-m** *arg1 arg2 ...*]

---

## ARGUMENTS

*file* is assumed to be in the format of /etc/passwd. An entry in *file* is changed only if the password presented by passwd matches the encrypted password in /etc/passwd.

**-m** performs a ypmake after *file* is modified. Any arguments following the option are passed to ypmake.

## DESCRIPTION

This daemon is not run by default nor started from inetd. To enable remote password updating for the network information service, put an entry for /etc/yppasswdd in the /etc/rc.local file of the host serving as the master for the NIS passwd map.

## EXAMPLE

If the NIS password file is stored as /etc/yp/src/passwd, to propagate password changes immediately, invoke the server as:

```
# /usr/etc/rpc.yppasswdd /etc/yp/src/passwd -m passwd DIR=/etc/yp/src
```

In this case, src is the NIS domain name.

## FILES

/etc/yp/Makefile

## SEE ALSO

ypmake
*Concentrix-2800 Commands and Applications:* yppasswd, make
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: passwd, ypfiles

# yppoll

## network information service get map version

Queries an NIS server (master or slave) for information about one map. The information returned is whether the map exists on the host, what version (order) number is in the map, and which server is the map's master.

---

**/usr/etc/yp/yppoll [-h *hostname*] [-d *domainname*] *mapname***

---

### ARGUMENTS

**-h** specifies the hostname of the NIS server. The default is the local host.

**-d** specifies the name of the NIS domain. The default is the domain for the local host.

*mapname* specifies the name of a map. Execute the command ypcat -k ypmaps to get a list of map names.

### DESCRIPTION

If the server is old (version 1), yppoll uses the old protocol to communicate, and also uses the old diagnostics in case of failure.

### SEE ALSO

rpcinfo, ypwhich ypserv
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# yppull

## network information service get map from master

Obtains for a slave NIS server a more recent copy of an NIS map, from the map's master if possible. If the host is the map's master, the command succeeds but no real action occurs.

---

**/etc/yp/yppull [-h** *hostname*] [**-d** *domainname*] *mapname*

---

### ARGUMENTS

**-h** specifies the hostname of the NIS server obtaining the map. The default is the local host.

**-d** specifies the name of the NIS domain. The default is the domain for the local host.

*mapname* specifies the name of a map. Execute the command ypcat -k ypmaps to get a list of map names.

### SEE ALSO

rpcinfo, yppush, yppoll, ypserv
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# yppush

## force propagation of a changed NIS map

Sends a copy of an NIS map from the master server to all other servers in the domain. If the host is not the map's master, the command succeeds but no real action occurs.

---

**/usr/etc/yp/yppush [-d *domainname*] [-v] *mapname***

---

### DESCRIPTION

Yppush is normally run only on the master NIS server by the Makefile in /usr/etc/yp after the master data bases are changed.

Yppush first constructs a list of NIS server hosts by reading the NIS map ypservers within the domain. Keys within the map ypservers are the ASCII names of the machines on which the NIS servers run.

A "transfer map" request is sent to the NIS server at each host, along with the information needed by the transfer agent (the program that actually moves the map) to call back the yppush. When the attempt has completed (successfully or not), and the transfer agent has sent yppush a status message, the results may be printed to stdout. Messages are also printed when a transfer is not possible, for example, when the request message is undeliverable or when the timeout period on responses has expired.

### ARGUMENTS

**-d** specifies the name of the NIS domain. The default is the domain for the local host.

**-v** (verbose) causes messages to be printed when each server is called and for each response. By default, only error messages are printed.

*mapname* specifies the name of a map. Execute the command ypcat -k ypmaps to get a list of map names.

### NOTE

In the current version, the transfer agent is ypxfr, which is started by the ypserv program. If yppush detects that it is speaking to an old protocol server, it uses the old protocol. However, yppush cannot detect if or when an old map transfer is performed. Yppush prints a message saying that an *old-style* message was sent. The system administrator should check to see that the transfer actually took place.

### FILES

/etc/yp/*domainname*/ypservers.{dir, pag}
list of NIS servers for an NIS domain
/etc/yp

## SEE ALSO

rpcinfo, yppoll, yppull, ypserv, ypxfr
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles
*Concentrix Network File System Manual:* chapters on NIS protocols and administration

# ypserv

## network information service server process

Looks up information in the NIS maps for the current domain. The ypserv daemon must be running on at least one host in an NIS domain. The host must be set up as an NIS server; see ypinit. Do not run /etc/ypserv on client hosts.

The ypserv daemon is typically invoked by /etc/rc.local. Also make sure /etc/ypbind is set up; see ypbind.

---

**/usr/etc/ypserv [-i] [-d] [-v]**

---

### ARGUMENTS

**-d** Debug. Causes ypserv to generate additional information on the usage of the Domain Name Service (DNS) to resolve host names.

**-i** Internet. Causes ypserv to attempt to use the DNS for name resolution when the NIS database resolution fails on host name lookup. This same effect can be caused without the -i option by using the -b switch to makedbm when creating the hosts database.

**-v** Verbose. Causes ypserv to not detach from the terminal and print dbm debugging info.

### DESCRIPTION

The network information service (NIS) facility is a simple network lookup service consisting of data bases and processes. The data bases are dbm files in a directory tree rooted at /etc/yp; they are described in ypfiles. The main processes are /etc/ypserv, the NIS data base server, and /etc/ypbind, the NIS binder.

The ypserv daemon's primary function is to look up information in its local database of NIS maps. The operations performed by ypserv are defined for the programmer by the header file <rpcsvc/yp_prot.h>. Communication to and from ypserv is by means of RPC calls. Lookup functions are described in the library routine ypclnt. Four lookup functions exist, all of which are performed on a specified map within some NIS domain: Match, Get_first, Get_next, and Get_all. The Match operation takes a key and returns the associated value. The Get_first operation returns the first key-value pair from the map, and Get_next can be used to enumerate the remainder. Get_all ships the entire map to the requester as the response to a single RPC request.

Two other functions supply information about the map, rather than map entries: Get_order_number and Get_master_name. In fact, both order number and master name exist in the map as key-value pairs, but the server will not return either through the normal lookup functions. (If you examine the map with makedbm, however, they are visible.) Other functions are used within the NIS subsystem itself, and are not of general interest to NIS clients. They include Do_you_serve_this_domain?, Transfer_map, and Reinitialize_internal_state.

## FILES

If the file /usr/etc/yp/ypserv.log exists when ypserv starts, error information is written to this file.

## SEE ALSO

makedbm, ypbind, ypinit, ypmake, yppoll, yppush, ypset, ypwhich, ypxfr
*Concentrix-2800 Commands and Applications:* domainname, pcat, ypmatch
*Concentrix-2800 System Reference*: *System Calls and C Library Routines*: dbm, ypclnt
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# ypset

## point ypbind at a particular server

Tells ypbind to get NIS services for the specified domain from the ypserv process running on *server*. Whether *server* is down or is not running ypserv is not discovered until an NIS client process tries to get a binding for the domain. At this point, the binding set by ypset is tested by ypbind. If the binding is invalid, ypbind attempts to rebind for the same domain.

---

**/usr/etc/yp/ypset [-V1|-V2] [-d *domain*] [-h *host*] *server***

---

### ARGUMENTS

**-V1** binds server for the (old) v.1 NIS protocol.

**-V2** binds server for the (current) v.2 NIS protocol. If no version is supplied, ypset first attempts to set the domain for the (current) v.2 protocol. If this attempt fails, ypset then attempts to set the domain for the (old) v.1 protocol.

**-d** *domain* uses domain instead of the default domain.

**-h** *host* sets ypbind's binding on *host* instead of locally. *Host* can be specified as a name or as an IP address.

*server* indicates the NIS server to bind to, and can be specified as a name or an IP address. If specified as a name, ypset attempts to use NIS services to resolve the name to an IP address, which works only if the node has a current valid binding for the domain in question. In most cases, *server* should be specified as an IP address.

### DESCRIPTION

Ypset is useful for binding a client node that is not on a broadcast net or is on a broadcast net that is not running an NIS server host. It also is useful for debugging NIS client applications, for example, where a NIS map only exists at a single NIS server host.

In cases where several hosts on the local net are supplying NIS services, ypbind can rebind to another host even as you attempt to find out if the ypset operation succeeded. That is, you can type ypset host1, type ypwhich, and get the reply host2, which can be confusing. This phenomenon is a function of the NIS subsystem's attempt to load balance among the available NIS servers, and occurs when host1 does not respond to ypbind because it is not running ypserv (or is overloaded), and host2, running ypserv, gets the binding.

### SEE ALSO

ypserv
*Concentrix-2800 Commands and Applications:* ypwhich
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# ypupdated

## server for changing NIS information

ypupdated is a daemon that updates information in the NIS, normally started up by inetd. ypupdated consults the file updaters in the directory /etc/yp to determine which NIS maps should be updated and how to change them.

By default, the daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure).

---

## /usr/etc/rpc.ypupdated [-is]

---

### ARGUMENTS

-i also accept RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of NIS maps in all networks.

-s accept only calls authenticated using the secure RPC mechanism (AUTH_DES authentication). This disables programmatic updating of NIS maps unless the network supports these calls.

### FILES

/etc/yp/updaters

### SEE ALSO

inetd, keyserv
*Concentrix-2800 Commands and Applications:* ypwhich
*Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: ypupdaters

# ypwhich

## determine network information service master

Prints the NIS server for a host or prints the map masters in a domain.

---

**/usr/bin/ypwhich [-d *domain*] [-V1] [-V2] [*hostname*]**

**/usr/bin/ypwhich [-t *mapname*] [-d *domain*] -m [*mname*]**

**/usr/bin/ypwhich -x**

---

### ARGUMENTS

Without arguments, ypwhich prints the NIS server for the local host.

**-d** *domain* specifies the name of an NIS domain. The default is the current domain for the local host.

**-m** [*mname*]
finds the master NIS server for the specified map, or (if *mapname* is omitted) for all maps in a domain. No *hostname* can be specified with -m. *Mname* can be a map name or a nickname (abbreviated name) for a map name (as described in ypcat).

**-t** *mapname*
inhibits nickname translation; useful if a map name is identical to a nickname.

**-V1** finds which server is serving v.1 NIS protocol client processes.

**-V2** finds which server is serving v.2 NIS protocol client processes. By default, ypwhich tries to a v.2 NIS server; if no v.2 NIS server is bound, ypwhich tries to find a v.1 server. Since servers and clients are backwards compatible, the version is usually not a user concern.

**-x** displays the map nickname table, which lists the known nicknames and associated mapnames.

*hostname* specifies the host for which the information is being supplied. Defaults to the local host.

### SEE ALSO

rpcinfo, yppush, ypserv, ypset
*Concentrix-2800 Commands and Applications:* ypcat
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles

# ypxfr

## transfer an NIS map

Moves an NIS map to the local host by making use of normal NIS services. Ypxfr creates a temporary map in the directory /etc/yp/domain (which must already exist), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. Ypxfr then deletes any old versions of the map and moves the temporary map to the real map name.

---

**/usr/etc/yp/ypxfr** [**-f**] [**-c**][**-h** *host*] [**-d** *domain*] [**-s** *domain*] [**-C** *t n i p*] [**-S**] *mapname*

---

### ARGUMENTS

**-c** does not send a Clear current map request to the local ypserv process. Use this flag if ypserv is not running locally at the time you are running ypxfr. Otherwise, ypxfr complains that it cannot talk to the local ypserv, and the transfer fails.

**-C** *t n i p* (only for use by ypserv) specifies that ypxfr should call back a yppush process at the host with IP address *i*, registered as program number *n*, listening on port *p*, and waiting for a response to transaction id *t*.

**-d** *domain* Specifies a domain other than the default domain.

**-f** forces the transfer to occur even if the version at the master is not more recent than the local version.

**-h** *host* gets the map from *host* regardless of what the map says the master is. If *host* is not specified, ypxfr asks the NIS service for the name of the master, and tries to get the map from there. *Host* can be a name or an Internet address in the form a.b.c.d.

**-s** *domain* Specify a source domain from which to transfer a map that should be the same across domains (such as the services.byname map).

**-S** This option causes ypxfr to require that the ypserv server, from which it obtains the maps to be transferred, is using "privileged" IP ports. Since only superuser processes are typically allowed to use privileged ports, this feature adds an extra measure of security to the transfer. If the map being transferred is a secure map, ypxfr sets the permissions on the map to 0600.

### DESCRIPTION

If ypxfr is run interactively, it writes its output to the terminal. However, if it is invoked without a controlling terminal, and if the log file /etc/yp/ypxfr.log exists, it will append all its output to that file. Since ypxfr is most often run from /usr/lib/crontab, or by ypserv, you can use the log file to retain a record of what was attempted, and what the results were.

For consistency between servers, ypxfr should be run periodically for every map in the NIS data base. Different maps change at different rates: the services.byname map may not change for months at a time, for instance, and may therefore be checked only once a day in the wee hours. You may know that mail.aliases or hosts.byname changes several times per day. In such a case, you may want to check hourly for updates.

A crontab entry can be used to perform periodic updates automatically. Rather than having a separate crontab entry for each map, you can group comands to update several maps in a shell script. Examples (mnemonically named) are in /etc/yp: ypxfr_1perday.sh, ypxfr_2perday.sh, and ypxfr_1perhour.sh. They can serve as reasonable first cuts.

## FILES

| | |
|---|---|
| /etc/yp/ypxfr.log | NIS transfer log file |
| /etc/yp/ypxfr_1perday | script to run one transfer per day, for use with cron |
| /etc/yp/ypxfr_2perday | script to run two transfers per day |
| /etc/yp/ypxfr_1perhour | script for hourly transfers of volatile maps |
| /etc/yp/*domain* | NIS domain |

## SEE ALSO

cron, yppush, ypserv
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: ypfiles, crontab
*Concentrix Network File System Manual:* chapters on NIS protocols and administration

# zic

## timezone compiler

The zic command reads text from the file(s) named on the command line and creates the time conversion information files specified in this input. If a *filename* is ' -- ' the standard input is read.

---

**/etc/zic** [ **-v** ] [ **-d** *directory* ] [ **-l** *localtime* ] [ **-p** *posixrules*] [ **-L** *leapsecondfilename* ] [ **-s** ] [ *filename ...* ]

---

### ARGUMENTS

**-v** complains if a year that appears in a data file is outside the range of years representable by the system call time values.

**-d** *directory*
creates time conversion information files in the named directory rather than in the standard directory named below.

**-l** *localtime*
uses the given time zone as local time. zic will act as if the input contained a link line of the form

```
link  timezone  localtime
```

**-p** *posixrules*
uses the given time zone's rules when handling POSIX-format time zone environment variables. zic will act as if the input contained a link line of the form

```
link  timezone  posixrules
```

**-L** *leap...* reads leap second information from the file with the given name. If this option is not used, no leap second information appears in output files.

**-s** limits time values stored in output files to values that are the same whether they're taken to be signed or unsigned. You can use this option to generate SVVS-compatible files.

### DESCRIPTION

Input lines are made up of fields. Fields are separated from one another by any number of white space characters. Leading and trailing white space on input lines is ignored. An unquoted sharp character (#) in the input introduces a comment which extends to the end of the line the sharp character appears on. White space characters and sharp characters may be enclosed in double quotes (" ") if they're to be used as part of a field. Any line that is blank (after comment stripping) is ignored. Non-blank lines are expected to be of one of three types: rule lines, zone lines, and link lines.

### Rule Lines

A rule line has the following form:

```
Rule  NAME  FROM  TO    TYPE  IN   ON       AT    SAVE  LETTER/S
```

For example:

```
Rule  USA   1969  1973  -     Apr  lastSun  2:00  1:00  D
```

The fields that make up a rule line are:

NAME — Gives the (arbitrary) name of the set of rules this rule is part of.

FROM — Gives the first year in which the rule applies. The word *minimum* (or an abbreviation) means the minimum year with a representable time value. The word *maximum* (or an abbreviation) means the maximum year with a representable time value.

TO — Gives the final year in which the rule applies. In addition to minimum and maximum (as above), the word *only* (or an abbreviation) may be used to repeat the value of the FROM field.

TYPE — Gives the type of year in which the rule applies. If TYPE is' –' then the rule applies in all years between FROM and TO inclusive; if TYPE is *uspres* the rule applies in U.S. Presidential election years; if TYPE is *nonpres* the rule applies in years other than U.S.Presidential election years. If TYPE is something else, then zic executes the command

```
yearistype year type
```

to check the type of a year: an exit status of zero is taken to mean that the year is of the given type; an exit status of one is taken to mean that the year is not of the given type.

IN — Names the month in which the rule takes effect. Month names may be abbreviated.

ON — Gives the day on which the rule takes effect. Recognized forms include:

| | |
|---|---|
| 5 | the fifth of the month |
| lastSun | the last Sunday in the month |
| lastMon | the last Monday in the month |
| Sun>=8 | first Sunday on or after the eighth |
| Sun<=25 | last Sunday on or before the 25th |

Names of days of the week may be abbreviated or spelled out in full. Note that there must be no spaces within the ON field.

AT Gives the time of day at which the rule takes effect. Recognized forms include:

| | |
|---|---|
| 2 | time in hours |
| 2:00 | time in hours and minutes |
| 15:00 | 24-hour format time (for times after noon) |
| 1:28:14 | time in hours, minutes, and seconds |

Any of these forms may be followed by the letter w if the given time is local "wall clock" time or s if the given time is local "standard" time; in the absence of w or s, wall clock time is assumed.

SAVE Gives the amount of time to be added to local standard time when the rule is in effect. This field has the same format as the AT field (although, of course, the w and s suffixes are not used).

LETTER/S
Gives the "variable part"(for example, the "S" or "D" in "EST" or "EDT") of time zone abbreviations to be used when this rule is in effect. If this field is ' – ' the variable part is null.

### Zone Lines

A zone line has the form:

```
Zone  NAME                    GMTOFF   RULES/SAVE   FORMAT   [UNTIL]
```

For example:

```
Zone  Australia/South-west      9:30     Aus          CST      1987 Mar 15 2:00
```

The fields that make up a zone line are:

NAME The name of the time zone. This is the name used in creating the time conversion information file for the zone.

GMTOFF
The amount of time to add to GMT to get standard time in this zone. This field has the same format as the AT and SAVE fields of rule lines; begin the field with a minus sign if time must be subtracted from GMT.

RULES/SAVE
The name of the rule(s) that apply in the time zone or, alternately, an amount of time to add to local standard time. If this field is ' –' then standard time always applies in the time zone.

FORMAT The format for time zone abbreviations in this time zone. The pair of characters %s is used to show where the "variable part" of the time zone abbreviation goes.

UNTIL The time at which the GMT offset or the rule(s) change for a location. It is specified as a year, a month, a day, and a time of day. If this is specified, the time zone information is generated from the given GMT offset and rule change until the time specified.

The next line must be a "continuation" line; this has the same form as a zone line except that the string "Zone" and the name are omitted, as the continuation line will place information starting at the time specified as the UNTIL field in the previous line in the file used by the previous line. Continuation lines may contain an UNTIL field, just as zone lines do, indicating that the next line is a further continuation.

### Link Lines

A link line has the form

```
Link  LINK-FROM    LINK-TO
```

For example:

```
Link  US/Eastern  EST5EDT
```

The LINK-FROM field should appear as the NAME field in some zone line; the LINK-TO field is used as an alternate name for that zone.

Except for continuation lines, lines may appear in any order in the input.

Lines in the file that describes leap seconds have the following form:

```
Leap  YEAR  MONTH  DAY  HH:MM:SS  CORR  R/S
```

For example:

```
Leap  1974  Dec      31      23:59:60  +      S
```

The YEAR MONTH DAY and HH:MM:SS fields tell when the leap second happened. The CORR field should be "+" if a second was added or "-" if a second was skipped. The R/S field should be (an abbreviation of) "Stationary" if the leap second time given by the other fields should be interpreted as GMT or (an abbreviation of) "Rolling" if the leap second time given by the other fields should be interpreted as local wall clock time.

For areas with more than two types of local time, you may need to use local standard time in the AT field of the earliest transition time's rule to ensure that the earliest transition time recorded in the compiled file is correct.

### FILES

| | |
|---|---|
| /etc/zoneinfo | standard directory used for created files |

### SEE ALSO

*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: tzfile

# APPENDIX B
# ADMINISTRATION PROCEDURES

This appendix describes the steps for performing a number of administration procedures.

> **Note:** These are procedures that *you* must step through; they are not automated scripts.

# Add_ethernet

## add system to Ethernet

Configure a system to communicate with other systems through Internet services over an Ethernet.

### PROCEDURE

Log in as root and do the following as necessary:

1. Physically connect an Ethernet transceiver to an Ethernet port on the back of the system or I/O expansion cabinet.

2. Type /usr/alliant/syscfg -d or examine a boot script to determine the names of your Ethernet controller(s). (Ethernet controller names are en0, en1 (CMC controller) or ex0, ex1 (Excelan controller), and so on.) If you have other than a single controller named en0 or ex0, adjust and/or add ifconfig lines (line 9) in /etc/rc.local.

   Invent a hostname for your system. The name must not contain spaces, tabs, or pound signs (#).

   The following example shows the Ethernet controller at ex0:

```
PATH=/etc:/usr/etc:/bin:/usr/ucb:/usr/bin:/usr/alliant

# myname is my symbolic name
# myhostid is my internet address in #.#.#.# format
# my-netmask is specified in /etc/networks
#
hostname nisi
#domainname mydomain
ifconfig ex0 -trailers `hostname`
ifconfig lo0 localhost
```

   If you adjust /etc/rc.local, reboot the system.

3. Determine the Ethernet address of the system from the most recent boot script. (Ethernet addresses are six hexadecimal digits separated by periods.)

   For a new network, invent a name for the network. The name must be 1–12 characters and must not contain spaces, tabs, or pound signs (#).

   Either obtain a registered Internet address or invent an address that conforms to Internet specifications. A class C (most typical) Internet address in decimal has four parts separated by periods. The first three parts specify the network and the last part specifies the host system. The first part must be 192–255 in value and the last three parts must be 0–255 in value.

4. Update the file /etc/networks on your system and on all other systems on your local networks. (If you are running network information service, update /etc/networks on the master server.) For each network, specify its name and its Internet address separated by spaces or tabs. The first entry must be for loopback-net, Internet address 127 (127.0.0.0). The following example demonstrates a system that accesses two local area networks:

```
# This file has only the essential entries for a user
# internet of 2 networks.
# The first entry must be present.
loopback-net    127             software-loopback-net
n1net           192.0.1         user-net-1
n2net           192.0.2         user-net-2
```

5. Update the file /etc/hosts on your system and on all other systems on your local networks. (If you are running network information service, update /etc/hosts on the master server.) For each host, specify its Internet address, its name, and any aliases, separated by spaces or tabs. The first entry must be for loopback-net, Internet address 127 (127.0.0.0). The following example demonstrates a system that accesses five local area networks (including itself):

```
127.0.0.1       localhost
192.0.1.1       n1s1    alpha1
192.0.1.2       n1s2    alpha2
192.0.1.3       n1s3    alpha3
192.0.2.1       n2s1    beta1
192.0.2.2       n2s2    beta2
```

6. Update the file /etc/ethers on your system and on all other systems on your local networks. (If you are running network information service and have a map based on /etc/ethers, update /etc/ethers on the master server.) For each host, specify its Ethernet address and its name (not an alias), separated by spaces or tabs. The following example demonstrates a system that accesses five local area networks (including itself):

```
8:0:20:0:A1:61 n1s1
8:0:20:0:A7:35 n1s2
8:0:20:0:C5:63 n1s3
8:0:20:0:C5:6D n2s1
8:0:20:0:C5:CB n2s2
```

7. Test the connections to your system and all other systems with the command /etc/ping (use CTRL C to stop ping):

```
# /etc/ping n1s2
```

   If the remote system is alive, you receive a series of messages from ping as packets are transmitted and acknowledged. When you type CTRL C, you receive a summary message announcing 0% packet loss.

   If the remote system is dead, you receive no packet acknowledgement messages. When you type CTRL C, you receive a summary message announcing 100% packet loss.

A bad response from your system or all systems indicates a problem on your system. Check the hardware connections; check the files /etc/networks, /etc/hosts, and /etc/ethers.

A bad response from one system indicates a problem on that system. If the system is up, check its hardware connections; check the files /etc/networks, /etc/hosts, and /etc/ethers.

8. Update the files /etc/hosts.equiv and /.rhosts on your system and on all other systems on your local networks, as desired. The file /etc/hosts.equiv permits automatic access to ordinary users from the named systems. The file /.rhosts permits automatic access to root users from the named systems. You can also add and update .rhosts files for individual users to extend their access. The .rhosts file must be owned by the user and must not allow w access to group and others. The following example demonstrates the contents of both /etc/hosts.equiv and /.rhosts to permit complete automatic access among five systems:

```
n1s1
n1s2
n1s3
n2s1
n2s2
```

9. If you need more than 64 pseudo terminal master/slave pairs for remote logins, add the necessary files with /dev/mkptynod. The following example adds another 16 pairs:

```
# /etc/mkptynod # 3
```

If you need more than 128 pseudo terminal pairs, you must first adjust the limit with the command /usr/alliant/ptynum.

Add the tty names to /etc/ttys as demonstrated in the following example:

```
ttyr0   none                    network
ttyr1   none                    network
                ...
ttyrf   none                    network
```

10. Ensure that all changes are propagated to other systems on your local area networks. If you are running the network information service (NIS), ensure that all changes are made on the master server(s) and that ypmake is executed.

## COMMANDS

| | |
|---|---|
| ifconfig | configure network interface parameters |
| mkptynod | make special files for pseudo terminals |
| ping | send icmp echo packet (network debugging) |
| ptynum | change nptys flag in system kernel |

## FILES

| | |
|---|---|
| /etc/ethers | Ethernet address data base |
| /etc/hosts | host name data base |
| /etc/hosts.equiv | list of trusted hosts |
| /etc/networks | network name data base |
| /etc/rc.local | site-specific start-up script |
| /etc/ttys | terminal initialization data |
| .rhosts | personal list of trusted hosts |

# Add_printer

## add printer to system

Make a parallel, serial, or remote printer available to the system. By default, one parallel printer is set up: the name is lp, the spool directory (print queue) is /usr/spool/lpd, the device is /dev/lp, which is a soft link to /dev/mlp0.

### PROCEDURE

Log in as root and do the following as necessary:

1. For a parallel printer, have the appropriate controller installed in the machine and physically connect the printer. For a serial printer, physically connect the printer to a tty port with a standard terminal-multiplexor cable; see "Terminal Devices" in Chapter 3.

2. For a parallel or serial printer, create a special file for the device. The conventional name for a parallel printer is mlp*c*, where *c* is the controller number. A printer always operates in character (raw) mode. The major number is 0x90 (144). The minor number is 0x*c*0001, where *c* is the controller number. The following example creates a special file for a printer connected to controller 1:

```
# /etc/mknod /dev/mlp1 c 0x90 0x10001
```

   The conventional name for a serial printer is /dev/tty*cu*, where *c* is the controller number and *u* is the unit number in hexadecimal (same as a terminal). A printer always operates in character (raw) mode. The major number is 0x98 (152). The minor number is 0x*c*000*u*, where *c* is the controller number (0x0 through 0x7) and *u* is the port number (0x0 through 0xf). Permissions should be 600 to prevent users from writing to the device. The following example creates a special file for a terminal connected to controller 0, unit f:

```
# /etc/mknod /dev/tty1f c 0x98 0xf
# /bin/chmod 600 /dev/tty1f
```

3. For a serial printer, update the file /etc/ttys to include an entry (a line) for the printer. The on flag must be absent or specified as off. The argument for terminal characteristics is usually 2 or std.9600 (direct-connect 9600-baud device), 6 or std.2400 (direct-connect 2400-baud device), or 7 or std.4800 (direct-connect 4800-baud device). The name of the file, without the directory prefix /dev/, must begin in column 1. The following example demonstrates an entry for /dev/tty0f:

```
tty0f   "/etc/getty std.9600" vt100   off
```

4. Make a spooling directory (print queue) under /usr/spool. The ownership should be daemon, the group daemon, and the permissions 775. The following example makes a spooling directory named lpq:

```
# cd /usr/spool
# mkdir lpq
# chown daemon lpq
# chgrp daemon lpq
# chmod 775 lpq
```

5. Edit /etc/printcap to add an entry for the new printer. The simplest technique is to copy the entry for lp, paste it, and modify it. The following example demonstrates an /etc/printcap entry for a printer named lpw:

```
lpw|local wide-paper line printer:\
        :lp=/dev/mlp1:sd=/usr/spool/lpw:lf=/usr/adm/lpw-errs:fo:\
        :rf=/usr/bin/fpr:of=/usr/lib/lpf:pw#132:
```

The entries for af (if used), lf, lp, and sd should be different for each entry in /etc/printcap. In the example, a printer named lpw queues its jobs in /usr/spool/lpw and prints to the device /dev/mlp1. The page width is 132.

For a serial printer, you must also specify the br entry (baud rate), and the fs and/or fc entry (set/clear sg flags; see the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: tty). In the following example, a printer named lpq queues its jobs in /usr/spool/lpq and prints to the device /dev/tty0f at a baud rate of 9600 with a flags setting of 06320 (CRMOD on, no parity, and XTABS on):

```
lpq|local letter quality serial printer:\
        :lp=/dev/tty0f:sd=/usr/spool/lpq:lf=/usr/adm/lpq-errs:fo:\
        :br#9600:fs#06320:\
        :rf=/usr/bin/fpr:of=/usr/lib/lpf:pw#80:
```

For a remote printer, the lp entry must be null, an rm entry must specify the name of the remote system, and an rp entry must specify the name of the printer on the remote system. Typically, only these entries and the sd entry are specified. In the following example, a printer named lpl queues its jobs in /usr/spool/lpl and prints to the printer named lpl on the system named farsys.

```
lpl|remote laser printer:\
        :lp:rm=farsys:rp=lpl:sd=/usr/spool/lpl:
```

## COMMANDS

| | |
|---|---|
| chgrp* | change file group |
| chmod* | change file access mode |
| chown | change owner |
| mknod | build a special file |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/gettytab | terminal configuration data base |
| /etc/printcap | printer capability data base |
| /etc/ttys | terminal initialization data |

# Add_quotas

## add file system quotas

Set quotas for selected users on one or more file systems. Enable quota checking on those file systems.

### PROCEDURE

Log in as root and do the following as necessary:

1. For each file system to employ quotas, change directory to the file system root and create an empty file named quotas. For example:

```
# cd /
# touch quotas
# cd /usr
# touch quotas
```

2. Set the time limit for stopping user activity in a file system after a soft limit is exceeded. The time limit is imposed on a per file system basis with edquota –t. You can omit this step if if the default of seven days is acceptable. For example:

```
# /etc/edquota -t
editing session
```

   Edit the file that appears, modifying the time limits for the various file systems. For example:

```
fs /usr blocks time limit = 1.00 day, files time limit = 0 (default)
fs / blocks time limit = 1.00 day, files time limit = 0 (default)
```

   See the administration command edquota for more detail.

3. For each user, set the quotas for each file system with edquota *username*. The quotas include soft and hard limits for blocks (aggregate space in bytes) and inodes (number of files). If quotas are the same for all users or groups of users, use the quotas for one user as a prototype for the other users. For example:

```
# /etc/edquota bob
editing session
# /etc/edquota -p bob judy jim jan
```

   For *editing session*, edit the file that appears, modifying the soft and hard limits. A limit that is left zero, or reset to zero, means no limit. For example:

```
fs /usr blocks (soft =     8M, hard =   10M) inodes (soft = 0, hard = 0)
fs / blocks (soft =    1M, hard =    2M) inodes (soft = 0, hard = 0)
```

   See the administration command edquota for more detail.

4. Add the quota option to the entries in /etc/fstab for which quotas are being maintained. For example:

```
/dev/xdv00r     /       4.2     rw,quota        1       1
/dev/xdv00f     /usr    4.2     rw,quota        1       2
```

5. Check and turn on quotas for all the file systems with quotacheck and quotaon. For example:

```
# /etc/quotacheck -v -a
# /etc/quotaon -v -a
```

   You can also apply these commands on a per file system basis. See the administration commands quotacheck and quotaon.

6. Uncomment the four lines in /etc/rc that check and enable file system quotas:

```
# Handle quotas (Remove leading #'s from each of the four lines
#                following the next blank line when activating
#                quotas for your system.)

                          echo -n 'checking quotas: '    >/dev/console
/etc/quotacheck -a -p                                    >/dev/console
2>&1
                          echo  'done.'                  >/dev/console
/etc/quotaon -a                                          >/dev/console
2>&1
```

## REMOVE QUOTAS

To disable quota checking, use the quotaoff command. To permanently disable quota checking, also remove the appropriate quota option in /etc/fstab. If you intend to never again enable quotas for a file system, you might remove its quotas file. If you permanently disable quotas for all file systems, you should also comment out the quotacheck and quotaon lines in /etc/rc.

## COMMANDS

| | |
|---|---|
| edquota | edit user quotas |
| quota* | display disk usage and limits |
| quotaon, quotaoff | turn file system quotas on and off |
| quotacheck | file system quota consistency checker |
| repquota | summarize quotas for a file system |
| rquotad | remote quota server |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| quotas | quotas file in root of each file system |
| /etc/fstab | static file system table |

# Add_terminal

## add terminal to system

Make an interactive terminal available to the system.

### PROCEDURE

Log in as root and do the following as necessary:

1. Physically connect the terminal or modem to a multiplexor port on the back of the system or I/O expansion cabinet. Note the controller number (0–3) and unit number (0–f). Ensure that the controller switches are set correctly and that the correct cables are used. (See *Concentrix-2800 System Administration*, Section 3.5.)

2. Create a special file for the device. The conventional name for a terminal is /dev/tty*cu*, where *c* is the controller number and *u* is the unit number in hexadecimal. A terminal always operates in character (raw) mode. The major number is 0x98 (152). The minor number is 0x*c*000*u* where *c* is the controller number and *u* is the unit number for terminals and outgoing modems, and 0x*c*008*u* for incoming modems. The following example creates a special file for a terminal connected to controller 1, unit f:

   ```
   # /etc/mknod /dev/tty1f c 0x98 0x1000f
   ```

   To make a set of special files for one controller (units 0x0 through 0xf), use the command /dev/mkttynod. The following example makes a set of special files for controller 1 (specify t for terminal files, m for incoming modems):

   ```
   # /dev/mkttynod xmv 1 t
   ```

   Ensure that the owner is root, and that permissions are 622 for terminals and incoming modems, and 666 for outgoing modems.

3. Update the file /etc/ttys to include an entry (a line) for the terminal. See "Ttys Format" below. The following example demonstrates an entry for /dev/tty1f as a direct terminal connection:

   ```
   tty1f "/etc/getty std.9600"    vt100 on secure
   ```

4. Enable the terminal by sending a hangup signal to the init process:

   ```
   # kill -HUP 1
   ```

## TTYS FORMAT

The format of /etc/ttys is as follows:

- **Terminal name (field 1)** – The terminal name starts in column 1. Use only the file name, not the full pathname.

- **Control program (field 2)** – Field 2 specifies a command line that invokes a program to initially control the terminal. The typical command is /etc/getty with one argument; the argument identifies a set of terminal characteristics described in the file /etc/gettytab (see below). The command line must be enclosed in quotation marks if it contains white space.

- **Terminal type (field 3)** – Field 3 specifies the type of terminal usually connected as the device. The terminal type must be a valid type specified in the file /etc/termcap (see the *Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: termcap).

The remaining fields specify flags and (optionally) a window control program:

- **On flag** – The keyword on activates the terminal for logins, that is, invokes the control program for the program. No on flag or the keyword off deactivates the terminal, that is, does not invoke the control program. Specify on for direct terminals and incoming modems unless you have a special reason to deactivate the connection. Specify off for outgoing modems and printers.

- **Secure flag** – The keyword secure permits the superuser (root) to log in on the terminal. Omit this flag to prevent superuser login from a terminal.

- **Window control program** – The keyword window= introduces a command line that invokes a windowing program that is executed prior to the control program. The command line must be enclosed in quotation marks.

The typical sets of terminal characteristics, taken from the file /etc/gettytab (see the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats*: gettytab) as distributed, are as follows:

- **0** – (aliases: d300, Dial-300) Dial-in line that cycles between 300, 1200, 150, and 110 baud when the user presses BREAK.

- **2** – (aliases: std.9600, 9600-baud) Direct-connect 9600-baud line.

- **3** – (aliases: D1200, Fast-Dial-1200) Dial-in line that cycles between 1200, 300, and 2400 baud when the user presses BREAK.

- **5** – (aliases: D300, Fast-Dial-300) Dial-in line that cycles between 300, 2400, and 1200 baud when the user presses BREAK.

- **h** – (aliases: h2400, High-Speed-Dialup-2400) Dial-in line that cycles between 2400, 1200, and 300 baud when the user presses BREAK.

- **m** – (aliases: m9600, High-Speed-Direct-9600) Direct-connect line that cycles between 9600, 2400, 1200, and 300 baud when the user presses BREAK.

## NOTES

- To configure the terminal at system installation, modify the script /dev/mknewdevs.

## COMMANDS

| | |
|---|---|
| kill* | terminate job or process |
| mknod | build a special file |
| mkttynod | make special files for terminals |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/gettytab | terminal configuration data base |
| /etc/term/*/* | format of compiled term file |
| /etc/terminfo | terminal capability data base |
| /etc/ttys | terminal initialization data |

# Add_user

## add a new user to the system

Update the appropriate files so that the system recognizes a new user. Create a home directory for the user and the necessary files to get started.

### PROCEDURE

Log in as root and do the following:

1. Decide on a login name (8 characters or less, and not beginning with a hyphen) and group(s) for the new account. Make sure that the login name does not already exist in /etc/passwd or /usr/lib/aliases. The group(s) can already exist or be new. The following example ensures that the name hacker is unique:

```
# cat /etc/passwd | grep hacker
# cat /usr/lib/aliases | grep hacker
#
```

   If you are using network information service, access the appropriate maps:

```
# ypcat passwd | grep hacker
# ypcat aliases | grep hacker
#
```

2. Using a standard editor, create new entries in the file /etc/group as necessary. See below for an example of group entries.

3. Using the command /etc/vipw, create an entry in the file /etc/passwd for the new user. By default, /etc/vipw places you in the vi editor; set the environment variable EDITOR to use a different editor (for example, setenv EDITOR emacs). Place the entry on one line as specified below in "Passwd Fields." Omit Field 2 (password); for Field 5 (GCOS), type the user's full name. The following example illustrates an entry in /etc/passwd:

```
hacker::301:300:Polly Hacker:/staff/usr/hacker:/bin/csh
```

   UID and GID numbers should start at 100; you should assign numbers to new users in ascending order.

4. (optional) (Adding the user to the primary group specified in step 3 is not necessary, but is recommended for documentation.) Using a standard editor, add the login name to existing entries in the file /etc/group. Each entry must be on one line as specified below in "Group Fields." Omit Field 2 (password). The following example illustrates modifications to entries in /etc/group to include hacker:

```
algorithms::300:able,baker,charlie,dog,echo,fox,gorn,hacker
vollyball::400:able,tango,walker,gorn,dog,hacker
```

5. Using a standard editor, add the login name to existing entries in the file /usr/lib/aliases and create new entries as desired. Each entry must start at the first position of a new line as specified below in "Aliases." The following example illustrates three entries in /usr/lib/aliases: the first and second entries are modifications to include the new user hacker; the third entry creates a new alias consisting of the new user's first initial and last name.

```
algorithms:     able, baker, charlie
                dog, echo, fox
                gorn, hacker
vollyball:      able, tango, walker
                gorn, dog, hacker
phacker:        hacker
```

6. Type the following command (unless step 5 is not used):

```
/usr/ucb/newaliases
```

   If you are adding more than one user, you may want to complete Steps 1 through 4 for each user, then perform Step 5 once.

7. Create (mkdir command) a login (home) directory for the user. Entry 6 in /etc/passwd must be the name of this directory.

```
# mkdir ~hacker
```

8. Copy (cp command) the skeleton files in /usr/skel to the user's login directory:

```
% cp /usr/skel/.??* ~hacker
```

9. Change the owner (chown command) and group (chgrp command) of this file to the login name of the new user. Use the primary group to which the user belongs. Optionally, place a soft link in /usr to access the home directory. The following example creates a login directory for hacker:

```
# chown -R hacker.algorithms ~hacker
# ln -s /staff/usr/hacker /usr/hacker
```

10. Create a password for the user (passwd command; the user can later change the password):

```
% passwd hacker
New password: <password not echoed>
Retype new password: <password not echoed>
```

11. Change the user's background information (passwd -f command). This command prompts for a name, location, office phone, and home phone, and inserts this information in the GCOS field of /etc/passwd.

```
% passwd -f hacker
Name []: Polly Hacker
Location []: 2C
Office Phone []: 123
Home Phone []: 617-263-8649
```

12. If necessary, use the edquota command to specify a limit to the user's usage of filesystems where quotas are in effect. To duplicate the quotas for the prototypical user of a particular filesystem, specify the -p option as follows (from the root directory of the given filesystem):

```
# /usr/etc/edquota -p username
```

13. Exit :

```
% exit
```

14. If you are using Network Information Service, login on the Network Information Service master for the passwd file and update the NIS data base:

```
# cd /etc/yp ; make
```

## PASSWD FIELDS

In the file /etc/passwd, place each entry on its own line (even if the line wraps on the terminal screen). Separate fields with colons; omit a field by typing two consecutive colons. The fields are as follows:

1. Login name. An alphanumeric string of 8 characters or less, not beginning with a hyphen. Login names must be unique on a system or on a local area network.

2. Password in encrypted form.

3. User identification number. An integer. User identification numbers must be unique on a system or on a local area network. Numbers 0, 1, 11, 66, and 900 are reserved. In general, UIDs should start at 100 and ascend in order.

4. Group identification number. An integer. Numbers 10, 11, and 31 are typically reserved for system accounts. In general, GIDs should start at 100 and ascend in order.

5. GCOS field. Contains additional information about the user.

6. Login (home) directory. The full pathname of the directory in which the user is placed upon login or upon execution of the cd command without an argument. Must match the directory created in Step 6.

7. Login shell (defaults to the Bourne shell /bin/sh). The typical shell on Concentrix-2800 is /bin/csh.

## GROUP FIELDS

In the file /etc/group, place each entry on its own line (even if the line wraps on the terminal screen). Separate the fields with colons; omit a field by typing two consecutive colons. The fields are as follows:

1. Group name. An alphanumeric string of 8 characters or less. Group names must be unique on a system or on a local area network.
2. Password in encrypted form.
3. Group identification number. An integer. Numbers 10, 11, and 31 are typically reserved for system accounts.
4. A list of users belonging to the group. Separate the names with commas.

A user can belong to no more than 16 groups. However, other UNIX systems such as Sun restrict a user to no more than 8 groups. If you are using network information service in conjunction with other systems, you must use the lesser value.

## ALIASES

In the file /usr/lib/aliases, start each entry at the first position of a new line. Continue an entry by starting a new line with white space (tabs or spaces). A pound sign (#) at the first position of a line means that the line is a comment. The fields are as follows:

1. Alias. An alphanumeric string. Aliases must be unique on a system or on a local area network. Terminate the alias with a colon.
2. Name(s). One or more login names or other aliases, separated by commas.

## COMMANDS

| | |
|---|---|
| cd* | change working directory |
| chgrp* | change file group |
| chown | change owner |
| edquota | set quotas for a user |
| exit* | exit from shell |
| mkdir* | create directory |
| passwd* | change login password |
| passwd -f | create finger information |
| newaliases | rebuild data base for mail aliases file |
| su* | substitute userid (superuser) |
| vipw | edit the password file |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| ~ | user login directory. |
| /etc/group | group file |
| /etc/passwd | password file |
| /usr/lib/aliases | aliases file for sendmail |
| /usr/lib/aliases.dir | binary representation of aliases file |
| /usr/lib/aliases.pag | binary representation of aliases file |
| /usr/skel/.??* | skeleton files |
| quotas | quotas file at root of given filesystem |

# Add_uucp

## add system to UUCP network

Configure a system to communicate with other systems through the UUCP (unix to unix copy) network. UUCP can communicate over modem connections, direct port-to-port connections, and local area networks (Ethernet).

### BASIC SETUP

Log in as root and do the following as necessary:

1. (Not necessary if you are using Ethernet.) Physically connect terminal ports on your system to incoming and outgoing (optional but recommended) modems; and/or connect terminal ports on your system to terminal ports on other systems. The /dev entry for an incoming modem must specify the minor number as the unit number plus 128; the permissions must be 666 (owner root). The /dev entry for an outgoing modem must specify the minor number as the unit number; the permissions must be 666 (owner root). Disable login processes on the port for an outgoing modem by removing the on option or specifying off in /etc/ttys for the port's entry.

2. Make the UUCP directories:

```
# cd /usr/src/usr.bin/uucp
# make mkdirs
# chmod 777 /usr/spool/uucp
```

3. Copy the configuration files from /usr/src/usr.bin/uucp/UUAIDS to /usr/lib/uucp:

```
# cd /usr/src/usr.bin/uucp/UUAIDS
# cp L* /usr/lib/uucp/
# cp USERFILE /usr/lib/uucp
# cd /usr/lib/uucp
# chown UUCP *
# chmod 400 /usr/lib/uucp/L.sys
```

   Edit the L files and USERFILE to reflect your configuration. The template files contain descriptions in the form of comments; L.sys contains the most extensive description. See also the *Concentrix-2800 System Reference*, "File Formats." For Ethernet connections, use the TCP line as a template.

4. (Ethernet connections only.) Near the end of the file /etc/rc, add the following line to the list of daemons started with /etc/rshd:

```
/etc/uucpd & echo -n ' ucpd'          > /dev/console
```

   Add the following line to the file /etc/services:

```
uucp          540/tcp              # uucp
```

5. Copy the maintenance scripts from /usr/src/usr.bin/uucp/UUAIDS:

```
# cd /usr/src/usr.bin/uucp/UUAIDS
# cp uu.daily /usr/local/lib/uucp.daily
# cp uu.hourly /usr/local/lib/uucp.6hours
# cd /usr/lib/uucp
# chown UUCP uucp.*
```

   In /usr/lib/crontab, uncomment the two lines that refer to UUCP scripts. You can copy the maintenance scripts to other locations and adjust /usr/lib/crontab accordingly. You can change the timing for uucp.6hours and rename that file.

   To poll a system that does not have an outgoing modem or that might not initiate contact often, call /usr/bin/uupoll from /usr/lib/crontab. The following example calls the system farsys at 00:15 and 12:15 each day. Schedule polling only in groups as large as the number of ACUs (automatic calling units) and do not overlap polling and the execution of uu.hourly.

```
15 0,12 * * * uupoll farsys
```

6. If you are using /usr/bin/tip, set up the files /etc/remote and /etc/phones. A template for /etc/remote is in /dist/etc/remote. Also see the *Concentrix-2800 System Reference: Miscellaneous Routines, Special Files, and File Formats.*

## ADD A REMOTE SYSTEM TO YOUR CONFIGURATION

To enable your configuration to recognize a remote system, obtain the system name and password from the administrator of the remote system, then log in as root and do the following:

1. Add an entry to /usr/lib/uucp/L.sys for the remote system. Make any adjustments to /usr/lib/uucp/L.aliases and /usr/lib/uucp/L-dialcodes that are necessary or desired.
2. Add an entry to /usr/lib/uucp/USERFILE for the remote system specified in /usr/lib/uucp/L.sys.
3. Add an entry to /etc/passwd for the login name specified in /usr/lib/uucp/USERFILE.

## ADD YOUR SYSTEM TO A REMOTE CONFIGURATION

Give the administrator of the remote system your system name and password. The remote system administrator must adjust /usr/lib/uucp/L.sys, /usr/lib/uucp/L.aliases (as necessary), /usr/lib/uucp/L-dialcodes (as necessary), /usr/lib/uucp/USERFILE, and /etc/passwd as shown above.

## COMMANDS

| | |
|---|---|
| cd* | change working directory |
| chmod* | change file access mode |
| chown | change owner |
| cp* | copy files, directories, and subtrees |
| cron | clock daemon |
| make* | maintain program groups |
| uupoll | poll a remote UUCP site |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/passwd | password file |
| /etc/remote | remote host description file |
| /usr/lib/crontab | clock daemon control file |
| /usr/lib/uucp/L.aliases | UUCP hostname alias file |
| /usr/lib/uucp/L.cmds | UUCP remote command permissions file |
| /usr/lib/uucp/L-devices | UUCP device description file |
| /usr/lib/uucp/L-dialcodes | UUCP phone number index file |
| /usr/lib/uucp/L.sys | UUCP remote host description file |
| /usr/lib/uucp/USERFILE | UUCP pathname permissions file |

# Add_yp

## add network information service (NIS) to network

Enable systems on a local network to access certain information through common data bases. Participating systems are of three types: master servers contain NIS maps and the base files for the maps; slave servers contain NIS maps but not the base files; clients access servers to obtain the common information. Different servers can be masters and slaves for different maps; this procedure uses the simple model of one master for all maps.

> **Note:** The term "network information service (NIS)" is now used to refer to the service formerly known as Sun Yellow Pages. The functionality remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

### CREATE NIS DATA BASE ON MASTER SERVER

Log in as root on the master server and do the following:

1. Type the following command:

   ```
   # /bin/domainname
   ```

   If the domain name is absent or incorrect, specify a domain name (inventing one if necessary). The following example specifies that the host system is part of a domain named d1:

   ```
   # /bin/domainname d1
   ```

2. Initialize NIS maps for data bases based on the following files: /etc/passwd; /etc/group; /etc/hosts; /etc/ethers; /etc/networks; /etc/services; /etc/rpc; /etc/protocols; /etc/netgroup; /usr/lib/aliases; /etc/usrclass.

   ```
   # /etc/yp/ypinit -m
   ```

   Use /etc/yp/makedbm to initialize other maps.

3. Start the server and binder daemons in that order:

   ```
   # /etc/ypserv
   # /etc/ypbind
   ```

4. Edit /etc/rc.local. Uncomment (remove the pound sign in column 1) the domainname line and change mydomain to the name of your domain. Uncomment the three lines for /etc/ypserv and the three lines for /etc/ypbind.

## CREATE NIS DATA BASE ON SLAVE SERVERS

Log in as root on the slave servers and do the following:

1. Type the following command:

   ```
   # /bin/domainname
   ```

   If the domain name is absent or incorrect, specify the domain name. The following example specifies that the host system is part of a domain named d1:

   ```
   # /bin/domainname d1
   ```

2. Copy the NIS maps from the master server. (<hostname> is the NIS server from which the maps should be read.)

   ```
   # /etc/yp/ypinit -s <hostname>
   ```

3. Start the server and binder daemons in that order:

   ```
   # /etc/ypserv
   # /etc/ypbind
   ```

4. Edit /etc/rc.local. Uncomment (remove the pound sign in column 1) the domainname line and change mydomain to the name of your domain. Uncomment the three lines for /etc/ypserv and the three lines for /etc/ypbind.

## CREATE NIS DATA BASE ON CLIENTS

Log in as root on the clients and do the following:

1. Type the following command:

   ```
   # /bin/domainname
   ```

   If the domain name is absent or incorrect, specify the domain name. The following example specifies that the host system is part of a domain named d1:

   ```
   # /bin/domainname d1
   ```

2. Start the binder daemon:

   ```
   # /etc/ypbind
   ```

3. Edit /etc/rc.local. Uncomment (remove the pound sign in column 1) the domainname line and change mydomain to the name of your domain. Uncomment the three lines for /etc/ypbind.

## UPDATE NIS DATA BASE

Log in as root on the master server and do the following:

1. Make the changes to the base files.
2. For changes made to the standard base files (/etc/passwd and so on), run the makefile in /etc/yp:

   ```
   # cd /etc/yp ; make
   ```

3. For changes made to non-standard base files, run /etc/yp/makedbm on the file and then run /etc/yp/yppush on the map.

## COMMANDS

| | |
|---|---|
| domainname | set or display name of current domain system |
| makedbm | make an NIS dbm file |
| ypcat | display values in an NIS data base* |
| ypbind | NIS binder process |
| ypinit | build and install NIS data base |
| ypmake | rebuild NIS data base |
| yppasswd | server for modifying NIS password file |
| yppoll | NIS get map version |
| yppull | NIS get map from master |
| yppush | NIS send map to slaves |
| ypserv | NIS server process |
| ypwhich | determine NIS master |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/passwd | password file |
| /etc/group | group file |
| /etc/hosts | host name data base |
| /etc/networks | network name data base |
| /etc/services | service name data base |
| /etc/protocols | protocol name data base |
| /etc/netgroup | list of network groups |
| /etc/yp/... | NIS data base and directory structure |

# Boot_system

## start system and bring to multi-user mode

Bring the system to the normal operating mode. Select one procedure depending on the current state of the machine; then perform follow-on actions as needed.

### Procedure A: FROM MULTI-USER MODE

If the system is in multi-user mode, log in as root, and type the following:

**/etc/shutdown -r** *time* [*message*]RETURN

*Time* can be specified as now for immediate shutdown, +*n* for shutdown in *n* minutes, or *hour:minute* for shutdown at a certain hour (24-hour clock) and minute. *Message* can be anything and provides a means for warning users of the impending shutdown. Examples:

```
# /etc/shutdown -r +5 to reconfigure - immediate rebootRETURN

# /etc/shutdown -r 13:00 to reconfigureRETURN

# /etc/shutdown -r nowRETURN
```

Specification of now on a system in use by multiple users is considered unfriendly.

Omission of the -r option to shutdown brings the system down to single-user mode; no reboot occurs. Substitution of the -H option halts the system placing you in console mode. Substitution of the -h option shuts the system down to Diagnostix-2800 (which terminates Concentrix-2800).

Substitution of the -B (stop in VMEbus booter), -D (stop in Diagnostix-2800 ddt), or -S (stop in Diagnostix-2800) option reboots the system and stops the reboot in the specified mode. You must respond properly from the console to continue. See "Follow-On Actions" below.

Substitution of the -b (stop in FX/2800 booter), -d (stop in Concentrix-2800 ddt), or -s (stop in single-user Concentrix-2800) option reboots only Concentrix-2800 (Diagnostix-2800 remains running) and stops the reboot in the specified mode. You must respond properly from the console to continue. See "Follow-On Actions" below.

The command reboot shuts down and reboots immediately without messages. The command halt shuts down to Diagnostix-2800 or console mode immediately without messages; no reboot occurs.

Rebooting with the file /fastboot in existence skips the file system checks normally performed during the reboot. The following commands create a /fastboot file: fastboot, fasthalt, and shutdown -f.

Use of reboot or halt on a system in use by multiple users is considered unfriendly. Use of fastboot from multi-user mode is not recommended.

**WARNING:** Never use fastboot after a panic or if you suspect any problem with the file system.

### Procedure B: FROM SINGLE-USER MODE

If the system is in single-user mode in Concentrix-2800 (usually as the result of a shutdown without the -r argument), type one of the following:

**/etc/reboot** RETURN

**/etc/fastboot** RETURN

The commands reboot and fastboot reboot immediately without messages; fastboot (or reboot with the file /fastboot in existence) skips the file system checks normally performed during the reboot.

Use of the -H option halts the system placing you in console mode. Use of the -h option shuts the system down to Diagnostix-2800 (which terminates Concentrix-2800).

Use of the -B (stop in VMEbus booter), -D (stop in Diagnostix-2800 ddt), or -S (stop in Diagnostix-2800) option reboots the system and stops the reboot in the specified mode. You must respond properly from the console to continue. See "Follow-On Actions" below.

Use of the -b (stop in FX/2800 booter), -d (stop in Concentrix-2800 ddt), or -s (stop in single-user Concentrix-2800) option reboots only Concentrix-2800 (Diagnostix-2800 remains running) and stops the reboot in the specified mode. You must respond properly from the console to continue. See "Follow-On Actions" below.

Examples:

```
# /etc/reboot RETURN

# /etc/reboot -H RETURN

# /etc/fastboot -d RETURN

# /etc/fastboot -s RETURN
```

The fastboot command (or reboot with the file /fastboot in existence) skips the file system checks normally performed during the reboot and can be used safely if you have shut down normally to single-user mode. The fastboot command can save considerable time, especially on a system with a large amount of secondary storage.

WARNING: Never use fastboot after a panic or if you suspect any problem with the file system.

Rather than rebooting, you can restart multi-user mode as shown below. Use this procedure after a normal shutdown when initialization of the system is not necessary.

1. Sync the disks (type sync RETURN), unmount all file systems (type /etc/umount -a RETURN), and sync the disks again.

2. Exit from the single-user shell (type CTRL d).

### Procedure C: FROM RESET

Use the following procedure if Diagnostix-2800 is hung. You should not use the reset procedure if only Concentrix-2800 is hung (see procedure E).

1. Set the front panel power switch to the power on position, *not* to the power locked on position (I/AUTO).
2. Press the reset switch on the front panel.
3. Turn the power switch to the power locked on position (I/AUTO) to proceed through console mode (reboot automatically). If you leave the power switch in the power on position, the machine remains in console mode (see "Follow-On Actions" below).

### Procedure D: FROM POWER OFF

Use the following procedure if the power is off or if the machine fails to boot from any other mode.

1. Ensure that the power is on for the console and printer.
2. If the power is on, turn the power switch on the front panel to the power off position.
3. Turn the power switch on the front panel to the power locked on position (I/AUTO) to proceed through console mode (reboot automatically). If you leave the power switch to the power on position, the machine remains in console mode (see "Follow-On Actions" below).

### Procedure E: FROM DIAGNOSTIX-2800

Use the following procedure if Concentrix-2800 is hung.

1. Connect to FX/2800 startup mode by typing CTRL ^ (hold down the control and shift keys, then press the 6 key) followed by a period.
2. Type y RETURN, or select other options, to restart Concentrix-2800. (Type ? RETURN to list the options.)

## FOLLOW-ON ACTIONS

The boot procedure may stop at various points, as listed below, depending on options used to reboot the system. You can also stop in console mode by placing the power switch in the power on position rather than the power locked on position. Otherwise, the boot procedure (as supplied by Alliant) runs to conclusion without further intervention. The boot procedure can take a few minutes, or it can take five to ten minutes or more, depending on whether a fastboot or reboot is in effect and on the amount of disk space that must be checked (for a reboot). The boot procedure is finished and the system is started in multi-user mode when Concentrix-2800 announces itself and types the login: prompt on the console.

### Console Mode (H mode)

If you stop in console mode, the following prompt appears:

```
cons>
```

The most common console commands are as follows:

- **options** – The options command displays and sets the flags that stop processing in other modes during system startup. Type o; the system fills out the word options, displays the current flags, and gives information on choices. A blank display indicates that no flags are set. Type any combination of BDSbds, followed by a carriage return, to set flags for the specified modes (type an invalid option to delete current flags; respond with ? to obtain a description of the flags). The following command sets the boot flags to stop in single-user mode, then resets them to stop in both Diagnostix-2800 and single-user modes:

  ```
  cons>Options = ; Enter RETURN to keep, SPACE to clear, or new options.
  New Options? S[RETURN]

  cons>Options = s; Enter RETURN to keep, SPACE to clear, or new options.
  New Options? Ss[RETURN]

  cons>
  ```

- **unix** – The unix command causes system startup to commence. Type u; the system fills out the word unix and processing continues. From a remote terminal, you must set parameters 6[RETURN] before typing unix.

  ```
  cons>unix
  ```

- **parameters** – The parameters command enables and disables the remote port. Type p; the system fills out the word parameters and a menu appears. Type 6 to enable the remote port and 7 to disable it.

At any time, type a question mark in response to the cons> prompt to obtain a list of console commands. The following example demonstrates the complete list of console commands:

```
cons>?

Commands are:
        B/W/L   - byte/word/long mode for examinations
        C       - enable IPC clock
        E/S     - open address in supervisor/selected function code space
        F       - set boot flags
        G       - transfer execution to address
        I/M     - initialize map to global/local memory
        P       - set parameters (console/remote baud rates, boot devices)
        R       - reset system
        T       - run tests
        U       - boot
        Dbcu    - device boot from specific unit (sets this unit as default)
        dbcu    - device boot from specific unit (does not change default)
        ?       - this list

An open location may be modified by entering a new value
An open location is terminated by: <cr> - close, <lf> - next, ^ - previous
A delete cancels the current command
An address ending with \ opens without read
All numbers are in hex
cons>
```

### VMEbus Booter Mode (B mode)

If you stop in VMEbus booter mode, prompts appear for an alternate kernel, debugger, and root device (/dev/xdv00p by default). Respond as follows:

RETURN — to specify the default. The defaults are the files dunix and ddt, on the first partition labeled as the Diagnostix-2800 root partition (dlabel command, -Dr*p* for Diagnostix-2800 root device). The search order for the first root partition is disk controllers and units in ascending order in the first VME chassis, then the second, and so on.

**xdv(***c***,***u***,***p***)***filename***** — to specify an alternate kernel image or debugger, where *c* is the controller number; *u* is the unit number; and *p* is the partition letter. For example: xdv(0,0,p)dunix or xdv(0,0,p)dddt. The disk specification can be omitted if the default is acceptable, for example, myunix to specify another kernel on the default device. The file name can be omitted if the default is acceptable, for example, xdv(1,0,p) to specify a kernel with the default file name on another device. Omit *filename* in response to the root device prompt.

**%h** — to display help text.

**%ls** — to list the files in the root partition of the current device.

**%r** — to read the ramdisk from a device.

ESCAPE — (applies only to dddt prompt) to skip loading a Diagnostix-2800 debugger.

The following example specifies an alternate kernel, and uses the defaults for the debugger and the root device:

```
Alliant VMEbus Booter [made: Wed May 9 19:32:23 EDT 1990]
Type %h for help

boot from:  xdv(0,0,p)dunix: mydunix[RETURN]
loaded @ 0x0: 469048.+63272.+50864.+38688.+39264. start @ 0x0
boot from:  xdv(0,0,p)dddt: [RETURN]
map_mem: start=670000, ddt_addr=3d0000, size=31, flags=a0000001
loaded @ 0x670000: 78856.+7064.+37496. start @ 0x670000
root device:  xdv(0,0,p): [RETURN]
```

Processing continues after your final response.

### Alliant ddt Mode (D mode)

The debugger does not prompt if you stop in Alliant ddt mode. A header line starting with the words Alliant ddt appears and processing stops. At this point, you can enter ddt commands; see the *Concentrix-2800 Systems Programming Guide*. To continue processing, type [ESCAPE]g or [ESCAPE]p (note: ddt echoes [ESCAPE] as a dollar sign). In the following example, the user stops in Alliant ddt mode and immediately continues:

```
Alliant ddt X.82 M cpu_id=0. [lnk: Wed Mar  1 16:03:53 EST 1989]
[ESCAPE]p
```

### Diagnostix-2800 Mode (S mode)

If you stop in Diagnostix-2800 mode during startup, or shut down or connect to Diagnostix-2800 mode, the prompt dunix-sh# or dunix(1)# appears. At this point, you can enter commands to the Diagnostix-2800 operating system. Type [CTRL]d to continue startup or reconnect to the previous mode, as shown below:

```
dunix-sh# [CTRL]d
```

After Concentrix-2800 is activated, you can connect to (obtain a shell in) Diagnostix-2800 mode from Alliant ddt 2800 mode, single-user mode, or Concentrix-2800 (multi-user) mode by typing [CTRL]^ (hold down the control and shift keys, then press the 6 key) followed by an exclamation mark, as demonstrated below:

```
Concentrix(118)# [CTRL]^!
dunix(1)#
```

### FX/2800 Startup Mode

If processing stops in Diagnostix-2800 mode during startup (S flag or –S option), processing also stops in FX/2800 startup mode; otherwise, processing proceeds. After you type CTRL d to exit Diagnostix-2800 (if you stop in this mode), you are prompted on how to proceed:

```
FX/2800 Startup

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

Respond yRETURN to continue system startup.

After Concentrix-2800 is activated, you can disconnect from Alliant ddt 2800 mode, single-user mode, or Concentrix-2800 (multi-user) mode by typing CTRL^ (hold down the control and shift keys, then press the 6 key) followed by a period. This action connects you to FX/2800 startup mode (without affecting the Concentrix-2800 mode), as demonstrated below:

```
Concentrix(118)# CTRL^.
[EOT]

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

To obtain a shell in Diagnostix-2800 mode, respond to the prompt with a lowercase h followed by a carriage return:

```
Reboot with flags set to "b"? (y for yes, ? for help): hRETURN
dunix(1)#
```

To terminate the Diagnostix-2800 shell and return to FX/2800 startup mode, type CTRL d in response to the dunix prompt:

```
dunix-sh# CTRL d
Reboot with flags set to ""? (y for yes, ? for help):
```

To reconnect to the previous mode, respond to the prompt with a lowercase t followed by a carriage return. Type another carriage return after you are connected.

```
Reboot with flags set to ""? (y for yes, ? for help): tRETURN
restarting tip
connected
RETURN
Concentrix(118)#
```

At any time in FX/2800 startup mode, you can respond to the prompt with a question mark to display the full list of options:

```
Reboot with flags set to ""? (y for yes, ? for help): ?

Response can be a combination of these characters:

 s - reboot concentrix, stop in single user concentrix
 b - reboot concentrix, stop in the concentrix booter
 d - reboot concentrix, stop in concentrix ddt
 h - "halt" (a dunix shell which, on exit, returns to the "Flags" prompt)

Or any combination of these characters:

 S - reboot dunix, stop in single user dunix
 B - reboot dunix, stop in the dunix booter
 D - reboot dunix, stop in dunix ddt
 H - halt dunix in IP cons> prompt

Or one of these single characters

' '  (space) reboot concentrix with no flags
 t - don't reboot, just restart tip
 y - reboot concentrix with shown flags
 R - enable / disable booting from ramdisk
 X - exit console_loop, go to single user Dunix (via 'kill -TERM 1')

Reboot with flags set to ""? (y for yes, ? for help):
```

## Diagnostic Monitor Mode

If the dm> prompt appears on the console, you are stopped in the diagnostic monitor. Stopping in the diagnostic monitor is unusual but might, for example, be caused by a modification to the file /etc/rc.diag. You can exit by typing the command exitRETURN.

### FX/2800 Booter Mode (b mode)

If you stop in FX/2800 booter mode, prompts appear for an alternate kernel, debugger, and console configuration file (/etc/console.config under Concentrix-2800 or /console/console.config under Diagnostix-2800, by default). Respond as follows:

RETURN — to specify the default. The defaults are the files unix and ddt, on the first partition labeled as the Concentrix-2800 root partition (dlabel command, -C*rp* for Concentrix-2800 root device). The search order for the first root partition is disk controllers and units in ascending order in the first VME chassis.

/dev/rxdv*cup*:/*filename* — to specify an alternate kernel image or debugger, where *c* is the controller number; *u* is the unit number; and *p* is the partition letter. For example: /dev/rxdv00r:/unix or /dev/rxdv00r:/ddt. Full names must be specified. Specify only *filename* in response to the cfg prompt.

! — to suspend the FX/2800 booter and obtain a shell in Diagnostix-2800 mode. Type CTRL d to terminate the Diagnostix-2800 shell and return to the FX/2800 booter.

The following example specifies an alternate kernel, and uses the defaults for the debugger and the console configuration file:

```
FX/2800 Booter

unix [/dev/rxdv00r:/unix]: /dev/rxdv00r:/myunixRETURN
ddt [/dev/rxdv00r:/ddt]: RETURN
cfg [console.config]: RETURN
```

Processing continues after your final response.

### Alliant ddt 2800 Mode (d mode)

The debugger does not prompt if you stop in Alliant ddt 2800 mode. A header line starting with the words Alliant ddt 2800 appears and processing stops. At this point, you can enter ddt commands; see the *Concentrix-2800 Systems Programming Guide*. To continue processing, type ESCAPE g or ESCAPE p (note: ddt echoes ESCAPE as a dollar sign). In the following example, the user stops in Alliant ddt 2800 mode and immediately continues:

```
Alliant ddt 2800 X.82 M cpu_id=0. [lnk: Wed Mar  1 16:03:53 EST 1989]
ESCAPEp
```

You can disconnect from Alliant ddt 2800 mode and connect to Diagnostix-2800 mode by typing CTRL ^! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from Alliant ddt 2800 mode and connect to FX/2800 startup mode by typing CTRL ^. (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

### Single-User Mode (s mode)

If you stop in single-user mode during startup, or shut down to single-user mode, the prompt Concentrix-sh# appears. At this point, you can enter commands to the Concentrix-2800 operating system. Type CTRL d to continue processing (that is, to start multi-user Concentrix-2800), as shown below:

```
Concentrix-sh# CTRL d
```

You can disconnect from single-user mode and connect to Diagnostix-2800 mode by typing CTRL ^ ! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from single-user mode and connect to FX/2800 startup mode by typing CTRL ^ . (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

### Concentrix-2800 Mode

The system is fully started when the login prompt appears on the console:

```
ALLIANT CONCENTRIX (mysystem)

login:
```

At this point, you can use the console as a login terminal.

You can disconnect from Concentrix-2800 mode and connect to Diagnostix-2800 mode by typing CTRL ^ ! (control-shift-6 followed by an exclamation mark); see above "Diagnostix-2800 Mode." You can disconnect from Concentrix-2800 mode and connect to FX/2800 startup mode by typing CTRL ^ . (control-shift-6 followed by a period); see above "FX/2800 Startup Mode."

### COMMANDS

| | |
|---|---|
| fsck | file system consistency check / interactive repair |
| getty | set terminal mode |
| fastboot | reboot system without checking disks |
| fasthalt | halt system in console mode without checking disks |
| halt | halt system in console mode |
| init | system initialization and process control |
| rc | command scripts for starting system |
| sync | update the super block |
| reboot | reboot system |
| shutdown | shut down system at a given time |
| umount | unmount file systems |

## FILES

| | |
|---|---|
| /dunix | Diagnostix-2800 kernel image |
| /etc/fstab | static information about filesystems |
| /etc/gettytab | terminal configuration data base |
| /etc/nologin | file whose existence prevents logins |
| /etc/rc | command script for starting Concentrix-2800 |
| /etc/rc.diag | command script for starting Diagnostix-2800 |
| /etc/rc.local | site-specific command script for starting Concentrix-2800 |
| /etc/ttys | terminal initialization data |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /usr/adm/utmp | login record of current users |
| /usr/adm/wtmp | login record of all logins and logouts |
| /vmunix | a symbolic link to /unix |

# Broadcast

## broadcast messages to users

Send messages to users on a system-wide basis.

### MESSAGE OF THE DAY

The login process displays the file /etc/motd on the user's terminal unless the user's login directory contains the file .hushlogin. You can maintain this plain text file with a standard editor. On a fairly open system, set the permissions on the file to read and write for all:

```
# chmod 666 /etc/motd
```

On a more restricted system, set the permissions on the file to read for all and write for root:

```
# chmod 644 /etc/motd
```

Since the file is displayed at each login, the typical practice is to keep the file short and restrict its use to critical messages.

### SYSTEM MESSAGES

The msgs command displays numbered messages from the directory /usr/msgs on the user's terminal. By default, the display starts with the oldest message that the user executing the command has not read, as indicated by the file .msgsrc in the user's login directory. Typically, the command msgs -f (the -f option displays No new messages) is placed in each user's .login file so that the user receives new messages at each login. The user can control reading the messages through a number of interactive responses. See msgs in the *Concentrix-2800 Commands and Applications*.

Messages are placed in the /usr/msgs directory by sending mail to the user msgs or invoking msgs with the -s option. You must be super-user.

### WRITE TO ALL USERS

The wall command sends a message (using the write command) to all users logged onto the system. See wall in the *Concentrix-2800 Commands and Applications*. You should be super-user when you use this command. The following example demonstrates the wall command:

```
# wall[RETURN]
Food alert - leftovers in main conference room[RETURN]
[CTRL]D
```

### COMMANDS

| | |
|---|---|
| msgs | read system messages |
| wall | write to all users |

## FILES

| | |
|---|---|
| ~/.hushlogin | file that inhibits message of the day |
| /etc/motd | message of the day |
| /usr/msgs/* | messages data base |

# Build_kernel

## build custom kernel

Build a new Concentrix-2800 operating system kernel, typically after modifying files in /usr/src/sys/CONCENTRIX. You must exercise caution in modifying the kernel. Inappropriate code and improper build procedures can impair the operation of the system. Do not bypass the backup step.

Alliant now ships distributed sources for customunix building in the directory /usr/src/sys/CONCENTRIX/dist. Be sure you have included the latest distribution. See the *Concentrix-2800 Installation Guide and Release Notes*.

1. Copy the file /unix (the system kernel) to a backup file. If you damage the kernel, you can boot from the saved kernel (by stopping in the booter and specifying the name of the saved kernel). Note the following example:

   ```
   # cp /unix /unix.29jul88
   ```

   Do not omit this step.

2. Move to the directory that contains the appropriate files.

   ```
   cd /usr/src/sys/CONCENTRIX
   ```

3. Make sure that the file called Makefile (in this directory) contains the names of any kernel source and object files that you added or modified.

   If you modified one of the sources distributed by Alliant, list the name of the object file in the DIST_OBJS macro. If you created a new source file, list the name of the object file in the USER_OBJS macro.

   In all cases, list the object file (and the method for compiling it from source) along with the other object file targets. Follow the pattern shown for other files.

   (Note that Makefile already contains the custom accounting routine.)

4. Execute make to build a kernel. During the build, any routines that you changed are compiled and linked into the kernel. The new kernel is named customunix. The command is as follows:

   ```
   make
   ```

5. Apply any custom patches (see Patch_kernel) that you normally apply to the kernel. Remember that (at this point) the kernel is named customunix.

6. Install the new kernel. After installation the kernel is named /unix. The command is as follows:

   ```
   make install
   ```

7. Reboot the system.

# Call_cs

## report problem to Alliant customer service

Report a problem to Alliant customer service and permit Alliant to diagnose the problem on your machine.

### TELEPHONE NUMBER (MASSACHUSETTS TOLL-FREE)

**1–800–453–5305**

### TELEPHONE NUMBER (U.S. OUTSIDE MASSACHUSETTS TOLL-FREE)

**1–800–225–4268**

### TELEPHONE NUMBER (OUTSIDE U.S. NOT FREE)

**617–486–4950**

### REQUIRED INFORMATION

Please have the following information ready when you place the call:

- Company name and site.
- System host name and serial number.
- Concentrix-2800 version number and version numbers of any other Alliant software associated with the problem.
- Information on any special software associated with the problem.
- Clear statement of the problem and the impact on your operations.
- Remote port telephone number.
- Root password.

The last two items permit Alliant to diagnose the problem on your machine. Of course, you can withhold this information if it jeopardizes the security of your machine. You may want to change the root password after Alliant leaves the machine.

# Connect_diag

## connect to Diagnostix-2800

Connect the console to Diagnostix-2800, which is running on the VMEDP. This procedure can be performed while the system is in Alliant ddt 2800 mode, single-user mode, or Concentrix-2800 (multi-user) mode.

### PROCEDURE

On the system console, do the following:

1. Type the escape character CTRL^ (hold down the control and shift keys, then press the 6 key) followed by a period. This action connects you to FX/2800 startup mode. In the following example, the user is logged in as root on a system that is in Concentrix-2800 mode:

```
Concentrix(118)# CTRL^.
[EOT]

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

2. Respond to the prompt with an uppercase X followed by a carriage return:

```
Reboot with flags set to "b"? (y for yes, ? for help): XRETURN
Going to single user Dunix (via kill -TERM 1)
dunix-sh#
```

To disconnect from Diagnostix-2800 and return to the starting mode, do the following:

1. Type CTRLd in response to a dunix prompt:

```
dunix-sh# CTRLd Thu Jul  5 10:42:13 EST 1990
FX/2800 Startup

Entering Console Boot Loop.

Reboot with flags set to ""? (y for yes, ? for help):
```

2. Respond to the prompt with a lowercase t followed by a carriage return. Type another carriage return after you are connected. Note the following example:

```
Reboot with flags set to ""? (y for yes, ? for help): tRETURN
restarting tip
connected
RETURN
Concentrix(118)#
```

# Disable_terminal

## disable (enable) interactive terminal

Make an interactive terminal unavailable (available) to the system.

### PROCEDURE

Log in as root and do the following as necessary:

1. Update the file /etc/ttys. Locate the entry for the terminal and remove the on flag or change it to off. (To enable a disabled terminal, add the on flag or change the off flag to on.) The following example demonstrates an entry that disables /dev/tty1f:

```
tty1f "/etc/getty std.9600"     vt100 off
```

2. Send a hangup signal to the init process:

```
# kill -HUP 1
```

### COMMANDS

| | |
|---|---|
| kill* | terminate job or process |

*See *Concentrix-2800 Commands and Applications* or man1.

### FILES

| | |
|---|---|
| /etc/ttys | terminal initialization data |

# Extend_swap

## extend area for swapping and paging

Adopt an alternate configuration for swapping and paging. The standard configuration uses the s partition on the root disk.

### PROCEDURE

The procedure is as follows. Exercise caution in using the /etc/dlabel command; ensure that the device name is correct before pressing RETURN.

1. Use the /etc/dlabel command to mark the disks that are being reconfigured. Use the -Cs option to specify a swap partition; omission of the -Cs option indicates that the disk does not contain a swap partition. The following examples assume a system with four Fujitsu M2344 disks mounted on controllers 0 and 1, units 0 and 1, with standard VMEbus names.

   Example 1: extend swap partition from s on root disk to s on root disk and s on disk 01.

   ```
   # /etc/dlabel -iw -Crr -Css -Drp -Dsq fj690s /dev/rxdv00w
   # /etc/dlabel -iw -Css fj690s /dev/rxdv01w
   ```

   Example 2: change swap partition from s on root disk to h on disk 01.

   ```
   # /etc/dlabel -iw -Crr -Drp -Dsq fj690s /dev/rxdv00w
   # /etc/dlabel -iw -Csh fj690s /dev/rxdv01w
   ```

   Example 3: extend swap partition from s on root disk to s on root disk and h on disk 10.

   ```
   # /etc/dlabel -iw -Crr -Css -Drp -Dsq fj690s /dev/rxdv00w
   # /etc/dlabel -iw -Csh fj690s /dev/rxdv10w
   ```

2. Use the /etc/dlabel command (without the -Cs option) to ensure that all other disks are not marked for swapping. (You can examine the status of the disk with the S option of /etc/diskf and leave it alone if it is correct.)

   Example 1: ensure that disks 10 and 11 do not have swap partitions.

   ```
   # /etc/dlabel -iw fj690s /dev/rxdv10w
   # /etc/dlabel -iw fj690s /dev/rxdv11w
   ```

   Example 2: ensure that disks 10, and 11 do not have swap partitions.

   ```
   # /etc/dlabel -iw fj690s /dev/rxdv10w
   # /etc/dlabel -iw fj690s /dev/rxdv11w
   ```

   Example 3: ensure that disks 01 and 11 do not have swap partitions.

   ```
   # /etc/dlabel -iw fj690s /dev/rxdv01w
   # /etc/dlabel -iw fj690s /dev/rxdv11w
   ```

3. If you are extending the swapping/paging area over multiple disks (examples 1 and 3), issue the /usr/alliant/extendswap command. Note that only one partition per disk can be marked for swapping and paging; at least one swap partition must be accessible to IP0; and all disks and controllers must be of the same type.

```
# /usr/alliant/extendswap on /unix
```

## COMMANDS

| | |
|---|---|
| diskf | disk utility program |
| dlabel | generated disk labels from prototype file |
| extendswap | enable/disable multiple-disk swapping/paging area |

## FILES

| | |
|---|---|
| /dev/* | device files |
| /etc/disktab | disk description file |

# Format_disk

## format and label disk

Format and label a disk. This procedure should not typically be needed since Alliant formats and labels disks before shipping them.

Alliant currently supports the following disk controllers:

- **Xylogics 754** (VMEbus) – Supports the SMD-E interface at a transfer rate of 3.0 MB/second.
- **Xylogics 7800** (VMEbus) – Supports the IPI-2 interface at a transfer rate up to 10.0MB/second.
- **Interphase V/SCSI 4210 Jaguar SCSI interface** – This device is a VMEbus SCSI host adaptor.

  **NOTE:** The following procedure assumes that you already have a Concentrix-2800 system running on hard disk. Formatting from a tape-based system is not supported for customer use.

The procedure shows specific disk type names, device names, and serial numbers. Substitute your own names and numbers as appropriate. Disk type names are as listed in the file /etc/disktab. Device names are as listed in /dev.

### Xylogics Controllers

The Xylogics 754 supports the SMD disk drive interface on a VMEbus. The following disk drives are used:

| | |
|---|---|
| fj690s | (fj_m2344) Fujitsu M2344 690MB 8" |
| fj1000s | (fj_m2382) Fujitsu M2382 1002MB 8" |

The Xylogics 7800 supports the IPI-2 disk drive interface on a VMEbus. The following disk drives are used:

| | |
|---|---|
| fj2600s | (fj_m2671P) Fujitsu M2671P 2.65GB 8" |

The Xylogics controllers support a concept known as slipping sectors that allows bad sectors to be mapped to alternate sectors on the same track. This scheme allows the bad sector to be replaced with almost no latency. Alliant disks as shipped from the factory have two (1GB drives) or four (fj690s and fj2600s) spare sectors on every track. While these are more sectors than most other computer vendors use, Alliant feels that the added safety makes it a good price/reliability tradeoff. Alliant also allows for alternate track mapping on the 1GB drives for tracks with more than two bad sectors.

The following procedure reformats and relabels a Fujitsu M2382 disk on controller 3, disk 1 (/dev/rxdv31z). Substitute "fj690s" for "fj1000s" to specify a Fujitsu M2344 disk.

**Note:** Ignore any "flaw overwrite prevented" messages.

1. Make disk device nodes:

```
# cd /dev
# mkdisknod vme 3 1
```

2. Read the internal flaw map from the disk and write it to a file. (If no internal flaw map exists, call Alliant Customer Service.)

```
# /etc/diskf /dev/rxdv31z B > /tmp/flaws31
```

3. Install a disk label in the operating system:

```
# /etc/dlabel -i fj1000s /dev/rxdv31z
```

4. Format and slip any bad sectors:

```
# /etc/diskf /dev/rxdv31z f /tmp/flaws31
```

5. Write the flaw file to the internal flaw map on the disk:

```
# /etc/diskf /dev/rxdv31z D /tmp/flaws31
```

6. Write the disk label to disk:

```
# /etc/dlabel -w fj1000s /dev/rxdv31z
```

### SCSI Controller

The Jaguar currently supports the Fujitsu M2266 1.2GB SCSI disk drive:

```
gen_scsi     (FUJM2266) Fujitsu M2266 1.2GB  5.25"
```

SCSI disks as delivered typically have a number of bad sectors; the disk manufacturer notes the bad sectors on a *hardware flaw map* on the disk. This is called the PLIST (primary defect list); it records information by cylinder, head, and sector. When you format a SCSI disk, the drive reads its PLIST, marks the bad spots, and formats around them.

If a sector on a SCSI disk goes bad as the drive is running, the drive recognizes the "grown" flaw and reports the block number(s) to the user in a console message. To fix the flaw, slip the sector with the s command of /etc/diskf, supplying the block number reported in the console message:

```
# /etc/diskf  /dev/rsd00w s <block_number>
```

The diskf utility then uses the "reassign blocks" command which fixes the flaw and records the bad sector in the drive's GLIST (the list of all "grown" flaws).

The drive slips bad blocks on *both* the PLIST and the supplied list during the formatting process, and saves the supplied list on the disk as the initial GLIST. Note that the diskf f command only formats *entire* SCSI disks.

The following procedure reformats and relabels a Fujitsu M2266 disk on controller 4, disk 1 (/dev/rsd31z).

1. Make disk device nodes:

```
# cd /dev
# mkdisknod sd 3 1
```

2. Read the internal flaw map from the disk and write it to a file. This saves the current GLIST.

```
# /etc/diskf /dev/rsd31z p > glist
```

3. Install a disk label in the operating system:

```
# /etc/dlabel -i gen_scsi /dev/rsd31z :bs=4k :nh=59 :ns=6 :sc=354
```

4. Format with the supplied list (and the PLIST) and slip any bad sectors:

```
# /etc/diskf /dev/rsd31z f glist
```

5. Write the disk label to disk:

```
# /etc/dlabel -w gen_scsi /dev/rsd31z :bs=4k :nh=59 :ns=6 :sc=354
```

## COMMANDS

| | |
|---|---|
| diskf | disk utility program |
| dlabel | generate disk labels from prototype file |

## FILES

| | |
|---|---|
| /dev/* | device files |
| /etc/disktab | disk description file |

## SEE ALSO

Label_disk
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: disktab

# Halt_system

## halt system in console mode

Halt the system in console mode.

### PROCEDURE

From multi-user Concentrix-2800, log in as root and type the following:

**/etc/shutdown –H** *time* [*message*]RETURN

*Time* can be specified as now for immediate shutdown, *+n* for shutdown in *n* minutes, or *hour:minute* for shutdown at a certain hour (24-hour clock) and minute. *Message* can be anything and provides a means for warning users of the impending shutdown. Examples:

```
# /etc/shutdown –H +30 to swap boards – 1 hourRETURN

# /etc/shutdown –H 13:00 to swap boards – 1 hourRETURN

# /etc/shutdown –H nowRETURN
```

Specification of now on a system in use by multiple users is considered unfriendly.

From console mode, you can shut off the power. You can also boot the system; see the follow-on actions of Boot_system.

From single-user Concentrix-2800, use one of the following:

```
# /etc/halt –HRETURN

# /etc/fasthalt –HRETURN

# /etc/reboot –HRETURN

# /etc/fastboot –HRETURN
```

### COMMANDS

| | |
|---|---|
| fastboot | reboot system without checking disks |
| reboot | reboot system |
| shutdown | shut down system at a given time |

### FILES

| | |
|---|---|
| /dunix | Diagnostix-2800 kernel image |
| /etc/nologin | file whose existence prevents logins |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /usr/adm/utmp | login record of current users |
| /usr/adm/wtmp | login record of all logins and logouts |
| /vmunix | a symbolic link to /unix |

# Install_batch

## install batch system

Set up a system for the submission of batch jobs.

### PROCEDURE

Log in as root and do the following:

1. Invoke an editor on the file /etc/qconf.

2. Modify the Batch Parameters section to reflect your system. Change the value of the hostname entry from localname to the actual name of your system. Optionally, change the value of the batch-queue entry (the default batch queue) from batch to some other name. The hostname change must be made (unless your system is actually named localname); the batch-queue value is typically left as is. The following example is for a system named s1:

```
#       @(#)etc.qconf  7.1     6/8/87
#       Batch Parameters
#
hostname        s1
sysmgr          root
batch-queue     batch
```

3. Modify the Batch Device Descriptions section to reflect the full complement of logical batch devices needed on your system. The distribution file provides entries for batch1, batch2, and net. You can supplement, change, delete, or comment out (pound sign in column 1) these entries. You should provide as many local devices as you want concurrently executing batch jobs; you should provide a network device if you want to submit jobs to other systems. Device and status entries for local and network devices must be as shown in the example. The following example names four local devices and one network device:

```
#       Batch Device Descriptions
#
# <dname>       <device>        <status>
batch1          /dev/null       anyform
batch2          /dev/null       anyform
batch3          /dev/null       anyform
batch4          /dev/null       anyform
net             /dev/null       skipmsg,anyform
```

4. Modify the Batch Queue Descriptions section to reflect the full complement of batch queues needed on your system. The distribution file provides entries for batch and batch@remote. You can supplement, change, delete, or comment out (pound sign in column 1) these entries. Leave the batch entry as is unless you changed the batch-queue entry above. Specify one remote entry for each system to which you want to send jobs; change remote in batch@remote to the actual name of the system. The following example names three local queues and two network queues:

```
#       Batch Queue Descriptions
#
# <qname>       <status>
batch
blittle
bbig
batch@farsys
batch@osys
```

5. Modify the Batch Mappings section to reflect the assignment of logical devices to batch queues on your system. The distribution file provides entries for batch and batch@remote. You can supplement, change, delete, or comment out (pound sign in column 1) these entries. Specify one or more entries for each local batch queue; specify one entry for each remote queue. Each entry for a local queue names a logical device that will execute jobs in the queue. The server for local queues must be /usr/lib/batch/shserver; the server for remote jobs must be /usr/lib/batch/netsend with two arguments that specify the name of a remote system and the name of a batch queue on that system. The following example maps four local devices to three local queues and a network device to two network queues:

```
#       Batch (Queue => Device) Mappings
#
# <qname>       <dname>         <server>
batch           batch1          /usr/lib/batch/shserver
batch           batch2          /usr/lib/batch/shserver
blittle         batch1          /usr/lib/batch/shserver
blittle         batch2          /usr/lib/batch/shserver
blittle         batch3          /usr/lib/batch/shserver
bbig            batch4          /usr/lib/batch/shserver
batch@farsys    net             /usr/lib/batch/netsend farsys batch
batch@osys      net             /usr/lib/batch/netsend osys batch
```

6. Exit from the editor. Be sure that the last characters of the file remain EOF:

```
# Do not remove this "EOF" --> EOF
```

7. Check the site-specific start-up script /etc/rc.local. The following lines must be present and not commented out:

```
if [ -f /usr/lib/batch/batchd ]; then
        rm -f /usr/spool/batch/qtmp/*
        /usr/lib/batch/batchd &
        (cd /usr/lib/batch; ./netdae -n 2 /usr/lib/batch/netrecv) &
        echo -n ' batch'                          >/dev/console
fi
```

   If you change /etc/rc.local, reboot the system or enter the commands manually.

8. Verify the existence of the directories /usr/spool/batch/qtmp and /usr/spool/batch/lock. If the directories do not exist, you must restore them from the distribution tape.

9. On each system that is a remote server, ensure that /etc/passwd (or the NIS password file) contains the following entry:

```
batch:*:2:1:batch:/:
```

## NOTES

- Systems to which you send remote jobs must also be running the batch facility.

## COMMANDS

| | |
|---|---|
| qdel* | delete batch request |
| qstat* | show batch request status |
| qsub* | submit batch request |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/qconf | batch configuration file |

# Install_mail

## install mail on a new system

Set up a new system for use of the mail and sendmail facilities.

### PROCEDURE

Log in as root and do the following:

1. Ensure that the host name of your system is correct by typing the following command:

   ```
   # /bin/hostname
   ```

   If the name displayed is incorrect (for example, if it is myname), edit line 7 of the start-up script /etc/rc.local to reflect the actual name of the system and ensure that the line is not commented out. You can reset the host name immediately as follows (assume that your host name is mars):

   ```
   # /bin/hostname mars
   ```

2. Set up paths so that the desired version of the mail command is invoked. The preferred version is /usr/ucb/mail, but the older /bin/mail is also available. A user can determine which version is in effect by typing:

   ```
   % which mail
   ```

   If the displayed pathname is not the desired version, the user should modify ~/.cshrc or ~/.profile to place /usr/ucb or /usr/bin first in the path command. The skeleton file /usr/skel/.cshrc places /usr/ucb first and you may want to adjust this file if /bin/mail is preferred on your system. Alternatively, the C shell user can place the command alias mail /usr/bin/mail in ~/.cshrc or ~/.login.

3. Create a system mail start-up file /usr/lib/Mail.rc if desired. Each user can add or override start-up options through the file ~/.mailrc. The file /usr/skel/.mailrc provides a template. See the *Concentrix System User's Guide*, "Mail Reference" for a list of options.

4. Add any system-wide aliases to the file /usr/lib/aliases. See "Aliases" below. Type the following command:

   ```
   # /usr/ucb/newaliases
   ```

5. Ensure that the three lines in /etc/rc.local that start the sendmail daemon are not commented out. You can start the daemon immediately by typing:

   ```
   # cd /usr/spool/mqueue
   # rm -f nf* lf*
   # /usr/lib/sendmail -bd -q1h
   ```

   The option -q1h means process the sendmail queue once an hour. You may want to change this option – see Appendix A: sendmail.

6. Copy the appropriate sendmail configuration file, /usr/lib/sendmail.main.cf (for use on a system that routes mail in a network of systems or that connects to UUCP for mail receiving and sending) or /usr/lib/sendmail.subsidiary.cf (for use in a local system environment or when the system just receives mail from other networked systems), to a work file, for example, /usr/lib/sendmail.custom. Make any desired edits, as indicated by comments in the file. Copy the work file to /usr/lib/sendmail.cf.

## ALIASES

In the file /usr/lib/aliases, start each entry at the first position of a new line. Continue an entry by starting a new line with white space (tabs or spaces). A pound sign (#) at the first position of a line means that the line is a comment. The fields are as follows:

1. Alias. An alphanumeric string. Aliases must be unique on a system or on a local area network. Terminate the alias with a colon.

2. Name(s). One or more login names or other aliases, separated by commas.

## COMMANDS

| | |
|---|---|
| hostname* | set/display host system name |
| /usr/ucb/mail* | send and receive mail |
| /bin/mail* | send and receive mail (old) |
| newaliases | rebuild data base for mail aliases file |
| sendmail | send mail over the internet |
| which* | show pathname for command |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/rc.local | site-specific command script for starting Concentrix-2800 |
| /usr/lib/aliases | aliases file for sendmail |
| /usr/lib/aliases.dir | binary representation of aliases file |
| /usr/lib/aliases.pag | binary representation of aliases file |
| /usr/skel/.??* | skeleton files |

# Label_disk

## label (relabel) disk

Label and relabel disks.

Disks are initially labeled as part of the formatting procedure; see Format_disk.

Disks are typically relabeled to remove unused partitions and to create non-standard partitions. Before relabeling a disk, back up its file systems. Exercise caution in modifying the disk description file /etc/disktab and in typing the /etc/dlabel command.

The disk description file /etc/disktab provides standard descriptions for the supported disks. The file is a plain ascii file that can be printed and edited. Each entry begins (after some comments) with the name of a disk type, then describes the disk and its logical partitions.

### PREVIOUS VERSION OF DISKTAB

The following example is the standard entry for a Fujitsu M2344. The :O entry specifies the starting cylinder of a partition; the :C entry specifies the size of a partition in cylinders. For example, the entries :Of#47 and :Cf#290 specify that partition f starts at cylinder 47 and is 290 cylinders in size. The cylinder size is the product of the :bt entry (bytes per track), :bs entry (block size), and :nh entry (number of heads). The last ten lines define b, d, r, f, s, p, q, h, w, and z partitions for the disk type.

```
#
# Fujitsu M2344 (690 Mb drive, 2.458MB/sec transfer rate, 8.0 inch drive)
#
# The following file storage and swap partitions are pre-defined:
#
#       b - boot partition for Concentrix
#       d - partition reserved for diagnostic reads and writes
#       r - root partition for Concentrix - 40mb
#       f - minimal user file partition for Concentrix - 256mb
#       s - swap for Concentrix - 182mb
#       p - root for Dunix - 60mb
#       q - swap for Dunix - 10mb
#       h - entire disk minus b and d partitions
#       w - entire disk minus the d partition
#       z - entire disk
#
fj690s|FJ690S|srm2344|Fujitsu M2344:\
        :ty=winchester:dt#11:ds=fj690s:\
        :vt#0:db#0:bt#40960:bs#512:\
        :nc#624:nh#27:ns#64:ac#0:ss#4:as#600:rm#3600:\
    :Ob#0:      :Cb#1:          \
    :Od#623:    :Cd#1:          \
    :Or#1:      :Cr#46:         \
    :Of#47:     :Cf#290:        \
    :Os#337:    :Cs#206:        \
    :Op#555:    :Cp#68:         \
    :Oq#543:    :Cq#12:         \
    :Oh#1:      :Ch#622:        \
    :Ow#0:      :Cw#623:        \
    :Oz#0:      :Cz#624:
```

The following example is the standard entry for a Fujitsu M2382:

```
#
# Fujitsu M2382 (1002 Mb drive, 3.0MB/sec transfer rate, 8.0 inch drive)
#
# The following file storage and swap partitions are pre-defined:
#
#       b - boot partition for Concentrix
#       d - partition reserved for diagnostic reads and writes
#       r - root partition for Concentrix - 43mb
#       s - swap for Concentrix - 315mb
#       p - root for Dunix - 60mb
#       q - swap for Dunix - 11mb
#       f - minimal user file partition for Concentrix - 390mb
#       h - entire disk minus b,d and t partitions
#       w - entire disk minus the d and t partition
#       z - entire disk minus the t partition
#       t - cylinder for spare tracks
#
fj1000s|FJ1000S|srm2382|Fujitsu M2382:\
        :ty=winchester:dt#19:ds=fj1000s:\
        :vt#0:db#0:bt#49728:bs#512:\
        :nc#745:nh#27:ns#80:ac#0:ss#2:as#604:rm#3600:\
    :Ob#0:      :Cb#1:          \
    :Od#743:    :Cd#1:          \
    :Or#1:      :Cr#39:         \
    :Of#40:     :Cf#353:        \
    :Os#393:    :Cs#285:        \
    :Op#688:    :Cp#55:         \
    :Oq#678:    :Cq#10:         \
    :Oh#1:      :Ch#742:        \
    :Ow#0:      :Cw#743:        \
    :Ot#744:    :Ct#1:          \
    :Oz#0:      :Cz#744:
```

## ENHANCED VERSION OF DISKTAB

Dlabel and its corresponding disktab have been enhanced to make it easier to create a disktab entry and to use that entry for more than one size drive. Changes in the disktab database or entries resulting from the addition of various types of disks and/or vendors are noted here. Note that sector size (bs) can now be 1 or 2KB; some SCSI disks support 4KB.

### Partition Table Fields

The :B, :F and :I fields are omitted from the partition tables. They are obsolete and SHOULD NOT BE SPECIFIED FOR SCSI DISKS. The fields available for specifying partition boundaries are as follows:

| | |
|---|---|
| :O | (origin) cylinder number of origin of partition |
| :S | (size) sizes (in blocks) of partitions |
| :C | (cylinders) sizes (in cylinders) of partitions |
| :U | (until) or "up to," specifies end of partition as being "up to but not including" the value. (default value: cylinders) |

A bare number specified in these fields is interpreted as some number of default units (for compatibility with previous versions of dlabel). Values may be given in other units by appending k (kilobytes), m (megabytes), g (gigabytes), c (cylinders) or b (blocks) to the number specified. In expressions, numbers without a suffix do not have an implicit multiplier; they are interpreted as byte counts.

All these values are eventually encoded as block numbers. The partition table in the header previously stored the partition origin as a 16-bit cylinder number. Since SCSI disks do not use any cylinder based size specifications, the origin is recorded in terms of a block number. Alliant's implementation uses the unused *iosize* field, a 32-bit value, to store the block number of a partition's origin. *If you do not omit the :I field*, it will overwrite the cylinder origin used by the SCSI disk driver.

### New Labelling Features

These new features may be used on SMD disks with the following restrictions:

1. Partition origins must still be a whole number of cylinders; partial values are rounded down to the nearest cylinder.

2. The ?.. and "end" variables cannot be used (yet).

You can use predefined variables and expressions with certain defined operators in the specification of partition values.

In the following discussion of the variables, P refers to a partition letter (a–z).

#### Variables

| | |
|---|---|
| startP | start of partition P |
| endP | end of partition P |
| end | end of the entire disk |
| sizeP | size of partition P |
| ?.. | automatically vendor-/drive-supplied information |

#### Operators

| | |
|---|---|
| +, –, *, / | usual meaning |
| && \|\| & \| ^ | usual meaning |
| <, >, <=,>= | usual |
| ==, != | usual meaning |
| e ? e ; e | same as C ?: operator, but using ';' since ':' is the disktab field delimiter |
| unary minus | usual meaning |
| () | usual meaning (grouping) |
| v roundupto b | rounds value v up to boundary b (b can be an expression or any of the suffix letters) |
| v rounddownto b | rounds value v down to boundary b |
| v roundto b | rounds value v to nearest boundary b |
| v? v; v | supply value |
| max(a,b) | the larger of a and b |
| min(a,b) | the smaller of a and b |

If circular references exist, the dlabel process is aborted.

The following example is the standard entry for a SCSI disk. Note that the = sign replaces the #. The :O entry specifies the starting cylinder of a partition; the :U entry specifies the end of a partition in cylinders. For example, the entries :Of=ends and :Ufstartp specify that partition f starts at the end of partition s and goes until the start of the p partition. In this example, size is specified only for the r and s partitions. The last ten lines define b, d, r, f, s, p, q, h, w, and z partitions for the disk type.

```
gen_scsi|general purpose scsi entry for drives around 1.2Gb:\
                :ty=scsi:ds=?:\
                :bs=?bs:\
                :nc=?nc:nh=?nh:ns=?ns:sc=?sc: \
        :Ob=0:                          :Ub=(256k+?bs) roundupto c: \:
        :Or=endb:                       :Sr=43m rounddownto c:      \
        :Os=endr:                       :Ss=max(end/4, 80m) rounddownto c: \
        :Of=ends:                       :Uf=startp:     \
        :Op=endp-60m rounddownto c:     :Up=startq:     \
        :Oq=endq-11m rounddownto c:     :Uq=startd:     \
        :Oh=endb:                       :Uh=startd:     \
        :Od=end-1c:                     :Ud=end:        \
        :Ow=0:                          :Uw=startd:     \
        :Oz=0:                          :Uz=startd:
```

The following list explains the entries for the partitions in the example shown above:

ds — normally specifies the drive name. If a ? is specified, the drive name is manufactured from the inquiry data (using 3 characters from the vendor name, and 5 characters from the drive name).

bs — specifies the block size. This parameter is required to be specified for an unformatted drive; Alliant specifies :bs=4k. If the drive has already been formatted, the variable ?bs contains the current block size in bytes.

nc,nh,ns — These fields are used by makefs, they are not looked at by the driver (or the drive). If you wish to set these fields to the drive's (nominal) parameters, the ?nc, ?nh, and ?ns variables contain them. For the Fujitsu M2266 SCSI disk drive, specify :ns=6 :nh=59.

sc — contains the number of sectors per cylinder. The diskf and dlabel commands read sc to determine where cylinders start and end. If the field is not specified, it will be filled in with (:nh *:ns). The variable ?sc contains the driver's idea of the number of sectors per cylinder. In particular, this gives the right value for Fujitsu drives; they have spare sectors on the last track. It should work correctly for other drives that have cylinder based spares. If the driver is not certain it can do the right thing, it returns (?nh * ?ns). For the Fujitsu M2266 SCSI disk drive, specify :sc=354.

The categories of cylinders, heads and sectors are useful for the file system optimizations attempted by makefs.

b — is the smallest number of cylinders which will hold 256KB + one block. This is to allow room for the booter (limited to 256KB) and the label (always takes block zero).

r — runs from the end of b for 43MB rounded down to the number of whole cylinders.

s starts from the end of r, its size is the larger of 1/4 of the whole disk or 80MB. The result is rounded down to the number of whole cylinders.

f runs from the end of s (which has a fixed relationship to the beginning of the disk), to the start of p (which has a fixed relationship to the end of the disk).

p runs from 60MB before its own end, rounded down to the number of whole cylinders, to the start of q.

q runs from 11MB before its own end, rounded down to the number of whole cylinders, to the start of d.

h runs from the end of b to the start of d.

d runs from one cylinder before the end of the media up until the end of the media.

w, z runs from the beginning of the disk up until the d partition.

### Notes on IPI-2 Disks

as On IPI-2 disks, this field is read from the drive and filled in automatically.

Do not modify an existing entry in /etc/disktab. To make an alternate entry for a disk type (using a text editor): copy the standard entry, paste it to the end of the file, change the name of the entry, and modify the contents as required.

Use the /etc/dlabel command to apply a new label to a disk.

### Adding User Partitions

To add user partitions, create a new entry in /etc/disktab and add the lines for the appropriate partitions. Then apply the label. The following partition names are available for users: g, i, j, k, l, n, o, u, and v.

For example, assume that you want to split the h (large user) partition on a Fujitsu M2344 disk into two equal partitions named u and v (so you can put two file systems on the disk).

#### Old Method

First do the necessary arithmetic to determine the number of sectors in each new partition and the starting cylinder of the second partition:

```
cylinders in h partition = :Ch#622 = 622
half of h partition in cylinders = 311
start of second partition = 311 + 1 = 312
```

The h partition is replaced by two non-standard partitions called u and v. Each new partition is 311 cylinders in size. The first partition starts at cylinder 1 (as does the h partition); the second partition starts at cylinder 312 (half the cylinders in the h partition plus 1). A new entry is created in /etc/disktab to define the u and v partitions, and remove unused partitions. (You might also remove the h partition.)

See the example below, where stars mark the lines where changes are made from the standard entry:

```
**fj690suv|FJ690S|srm2344|Custom Fujitsu M2344:\
  **:ty=winchester:dt#11:ds=fj690suv:\
    :vt#0:db#0:bt#40960:bs#512:\
    :nc#624:nh#27:ns#64:ac#0:ss#4:as#600:rm#3600:\
  :Ob#0:       :Cb#1:          \
  :Od#623:     :Cd#1:          \
  :Or#1:       :Cr#46:         \
  :Of#47:      :Cf#290:        \
  :Os#337:     :Cs#206:        \
  :Op#555:     :Cp#68:         \
  :Oq#543:     :Cq#12:         \
  :Oh#1:       :Ch#622:        \
**:Ou#1:       :Cu#311:        \
**:Ov#312      :Cv#311:        \
  :Ow#0:       :Cw#623:        \
  :Oz#0:       :Cz#624:
```

The entry is applied to the appropriate disk with the following command using standard names):

```
# /etc/dlabel -iw fj690suv /dev/rxdv10w
```

### New Method

To make o and u, corresponding to the first and second half of the h partition, respectively, add the following entries:

```
:Ou=endb:        :Su=sizeh/2:    \
:Ov=endu         :Uv=startd:     \
```

The entry is applied to the appropriate disk with the following command using standard names):

```
# /etc/dlabel -iw fj690suv /dev/rxdv10w
```

### Disk Stripes

To relabel a disk that currently participates in a disk stripe:

3. Disassemble the disk stripe with /etc/unmkstripe.
4. Relabel the constituent disks with /etc/dlabel.
5. Remake the disk stripe with /etc/mkstripe.

For example:

```
# /etc/unmkstripe /dev/rstripe0
# /etc/dlabel -iw fj690suser /dev/rxdv01w
# /etc/dlabel -iw fj690suser /dev/rxdv10w
# /etc/mkstripe /dev/rstripe0 /dev/rxdv01h /dev/rxdv10h
```

### COMMANDS

| | |
|---|---|
| diskf | disk utility program |
| dlabel | generate disk labels from prototype file |

### FILES

| | |
|---|---|
| /dev | device files |
| /etc/disktab | disk description file |

### SEE ALSO

Format_disk
*Concentrix-2800 System Reference*: *Miscellaneous Routines, Special Files, and File Formats*: disktab

# Make_file_system

## make and mount a new file system

Make a new file system, mount it, and update the file system table. Ensure that the disk partition(s) you choose for the new file system does not already contain a file system or overlap another partition that does. For striped file systems, all disks and controllers must be of the same type.

### PROCEDURE

Log in as root and do the following:

1. *For a striped file system only,* make special files (raw and block mode) in /dev for the striped device, if the files do not already exist. The conventional names are /dev/stripe*n* and /dev/rstripe*n*. An example follows:

   ```
   # /etc/mknod stripe0 b 200 0
   # /etc/mknod rstripe0 c 200 0
   ```

   Edit the file /etc/stripetab. (Create the file if it does not exist.) On a new line, specify the raw device name of the striped device followed by the raw device names of the partitions that will make up the striped file system. The names must be separated by spaces or tabs. A comment can be placed in the file by starting a line with a pound sign (#). The following example demonstrates a striped file system on the device /dev/rstripe0 consisting of the partitions /dev/rxdv01h and /dev/rxdv11h:

   ```
   # /bin/cat /etc/stripetab
   # file system for project x uses two disks on two controllers
   /dev/rstripe0  /dev/rxdv01h   /dev/rxdv11h
   ```

   Issue the following command:

   ```
   # /etc/mkstripe -a
   ```

2. Use the makefs command to create the file system. For 4.3BSD and track-allocated file systems, specify the raw device name of the partition that contains the file system. For striped files systems, specify the raw device name of the striped device. Specify a striped file system as type tfs. Note the following examples:

   ```
   # /etc/makefs -t 4.3 /dev/rxdv01h      4.3BSD

   # /etc/makefs -t tfs /dev/rxdv01h      track-allocated

   # /etc/makefs -t tfs /dev/rstripe0     striped
   ```

   Save the information displayed by the makefs command, preferably by writing it in a notebook. The alternate superblock numbers are of particular importance.

3. Make a directory file in the root file system to serve as the mount point for the new file system, for example:

   ```
   # /bin/mkdir /projectx
   ```

4. Mount the new file system. Use the block device name to specify the disk partition or striped device. Specify the mount point in the existing file system as the second parameter. For example:

```
# /etc/mount /dev/xdv01h /projectx      4.3BSD, track-allocated

# /etc/mount /dev/stripe0 /projectx     striped
```

5. Update the file system table /etc/fstab. The simplest way to effect the update is to recreate the file using the -p option of /etc/mount; save the old file first.

```
# /bin/cp /etc/fstab /etc/fstab.old
# /etc/mount -p > /etc/fstab
```

Alternatively (for example, if all file systems are not currently mounted), add a line to /etc/fstab for the new file system. Ensure that the pass number (last parameter) is greater than 1.

```
# /bin/cat >> /etc/fstab
/dev/xdv01h /projectx 4.3 rw 1 2      4.3BSD
/dev/xdv01h /projectx tfs rw 1 2      track-allocated
/dev/stripe0 /projectx tfs rw 1 2     striped
CTRL D
```

## COMMANDS

| | |
|---|---|
| chgrp* | change file group |
| chown | change owner |
| mkdir* | create directory |
| makefs | construct a file system |
| mknod | build a special file |
| mkstripe | make striped partition |
| tunefs | tune up an existing file system |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| /etc/fstab | static information about file system |
| /etc/stripetab | striped partition data base |

# Patch_kernel

## modify variables in the system kernel

Modify the system kernel to change functions and limits. You must exercise caution in patching the kernel. Selecting inappropriate values and mistyping values can impair the operation of the system. Do not bypass the backup and logging steps.

### PROCEDURE

Log in as root and do the following:

1. Copy the file /unix (the system kernel) to a backup file. If you damage the kernel, you can boot from the saved kernel (by stopping in the booter and specifying the name of the saved kernel). Note the following example:

   ```
   # cp /unix /unix.29jul90
   ```

   Do not omit this step.

2. Look up the variable under "Kernel Variables" below. Check values, limits, and sizes (word or longword).

3. Invoke the assembly-level debugger adb on the kernel for writing and for accessing kernel memory in the running system:

   ```
   # adb -w /unix /dev/kmem
   ready
   ```

   To fix the /unix file on the root disk, rather than memory in the running system, leave out /dev/kmem, as follows:

   ```
   # adb -w /unix
   ready
   ```

4. Examine the current value of the variable using one of the following adb commands:

   *variable-name*?D [for a longword; decimal]
   *variable-name*?d [for a word; decimal]
   *variable-name*?X [for a longword; hexadecimal]
   *variable-name*?x [for a word; hexadecimal]
   *variable-name*?s [for a string]

   Note the following example:

   ```
   nproc?D
   _nproc:         202
   nproc?X
   _nproc:         0xCA
   versionname?s
   _versionname:   Alliant Concentrix-2800 1.0.0
   ```

5. Change the value. Make sure the new value is within specifications and guidelines. Specify the value in hexadecimal as follows:

   *variable-name*?W *hex-value* [for a longword]
   *variable-name*?w *hex-value* [for a word]

   Verify the new value. Note the following examples. Use / to patch the disk and ? to patch running memory:

   This example patches the disk.

```
nproc/W 12F
_nproc:         0xCA            =       0x12F
nproc/D
_nproc:         303
```

   This example patches running memory:

```
nproc?W 12F
_nproc:         0xCA            =       0x12F
nproc?D
_nproc:         303
```

6. Exit from adb by typing $q or CTRL D.

7. Log the change. The recommended procedure is to maintain a shell script named /usr/patches/unix.sh (make the directory /usr/patches if it does not exist) with the following format:

```
#
adb -w /unix << EOF
adb write command
...
EOF
```

   The following unix.sh script logs one patch to nproc:

```
#
adb -w /unix << EOF
nproc?W 12F
EOF
```

   Do not omit this step. If you have a problem with your system, Alliant customer service must know about any changes to kernel variables. If you install a new kernel (for example, if you upgrade your system), you can run this script to apply your old tuning patches.

8. Reboot the system unless the variable is dynamic. (Most of them are not.)

## KERNEL VARIABLES

An abridged list of kernel variables follows. Do not change the ones marked information only. Where a kernel variable can be modified by a command, you should use the command to make the change (rather than patching).

**breaktop** (longword) virtual memory limit per process. Default is 0x20000000 (536MB). The maximum limit is the bottom of the process stack, which typically exceeds 0xC0000000 (3.2GB). However, the ABI specifies a limit of 0x80000000 (2GB), and programs designed to execute on other systems according to ABI specifications must not exceed this limit. Sufficient paging file space must exist to support the virtual size of a process.

**cacstack** (longword; information only) lower bound of cactus stack, which holds per-processor data during concurrent operations. Each processor has a separate cactus stack in physical memory; however, the virtual address is the same for all stacks. Use cacstacksegs to change the size of the cactus stack.

**cacstacksegs**
(longword) size of cactus stack in segments; each segment is exactly 4MB. The minimum is 1 segment and the maximum is 8 segments. Defaults to 8 (32MB).

**cacstacksize**
(longword; information only) size in bytes of each cactus stack (one per CE during concurrent operations). Use cacstacksegs to change the size of each cactus stack.

**desfree** (longword) number of free pages before paging/swapping starts; 0 lets system decide. Default is 0. Do not reset this value from 0 if you have less than 80 megabytes of real memory. If your real memory configuration changes and is equal to or greater than 80 megabytes, you can change desfree up to approximately 10% of real memory. Dynamic – system does not have to be rebooted.

**enbproflag** (longword) kernel profiling: 0 – disable; 1 – enable. Default is 0.

**ex_n_pages**
(word) network incoming packet space in half-pages (hpages); 0 lets the system decide. Default is 0. The command netstat -m lists the half-pages used/allocated.

**ISR_commsize**
(longword) amount of space allocated for the communications region for user-writtens interrupt service routines, specified in 4096-byte pages. Default is 1.

**ISR_loadsize**
(longword) amount of space allocated for the text and data for user-written interrupt service routines, specified in 4096-byte pages. Default is 2.

**mappedisable**
(longword) disable mapped files: 0 – off (enable); 1 – on (disable). Default is 0.

**MAX_SHARED_CREATE**
(longword) maximum temporary shared memory region in units of 4096 bytes. Default is 512.

**MAX_SHARED_PERMCREATE**
(longword) maximum permanent shared memory region in units of 4096 bytes. Only the superuser can exceed this limit; 0 prevents normal users from creating permanent shared memory regions. Default is 0.

**maxsharesize**
(longword) maximum memory for shared images in units of 4096 bytes. Default is 0 (approximately 75% of usermem). Use the administration command /usr/alliant/setmaxshare to change this variable. Do not exceed usermem.

**nbuf** (longword) buffers for cache I/O (number of buffers); 0 lets the system decide. Default is 0. The system (by default) allocates I/O buffers as follows: 10% of memory up to 128MB, plus 5% of memory over 128MB up to 1GB, plus 2% of memory over 1GB.

**nfile** (longword) memory allocated for file structures (number of files). Default is nproc plus 16 plus maximum users, times 32, divided by 10, plus 32, for example, 1990 for a 64-user system.

**nfs_wakeup_one_nfsd**
(longword) when set to 1, wakes up only the first process in the sleep queue for a socket when input occurs on the socket. This behavior improves the performance of NFS where typically many nfsd processes are on the sleep queue. However, if a program on your system depends on multiple processes being awakened, set this variable to 0. Default is 1.

**ninode** (longword) memory allocated for inodes (number of inodes). Default is nproc times 3, for example, 1596 for a 64-user system.

**nmbclusters**
(longword) cluster network buffers (mbufs) in pages (2 per page). These are allocated from space reserved by nmbufs. Default is zero, which becomes 50% of nmbufs. The command netstat -m lists the cluster mbufs used/allocated.

**nmbufs** (longword) network buffers in pages (16 per page); 0 lets the system decide. Default is 0. The command netstat -m lists the mbufs used/allocated.

**nproc** (longword) maximum number of processes; the recommended range is 102–1002. Default is maximum users times 8, plus 20, for example, 532 for a 64-user system.

**nptys** (longword) maximum pseudo tty ports. Default is 64. Use the administration command /usr/alliant/ptynum to change this variable.

**pagesharedflag**
(longword) when ON, pages that are parts of mapped files or of shared regions are eligible to be paged out. Default is ON. If you have a stable mix of jobs with long gaps between references, and plenty of memory, performance may be better with this flag OFF. If you have limited memory and/or a large number of shared regions (visualization programs), this flag should be ON.

**pallocsize** (longword) memory for system use in pages (takes away from usermem). Default ranges from 512 on small memory systems (32MB) through 2048 on large memory systems (256MB and up). Do not change without consulting Alliant customer service.

**rlimit_core** (longword) maximum size of core dump in bytes. Default is 20,971,520 (5120 pages).

**setclassdef** (longword) enables all users to set process class: 0 – all; 1 – superuser only. Default is 1. Use the administration command /usr/alliant/setclassdef to change this variable.

**soft_net_cpu**
(longword) IP doing Ethernet softnet processing. Default is 28 (machine 28, that is, IP0). Do not copy a /unix file with this change to another system (in general, a bad practice).

**stackbottom**
(longword; information only) lower bound of the user stack. The lower this value the larger the stack (the stack grows from higher to lower memory). Usrstack minus stackbottom equals the size of the user stack. Use usrstacksegs to change the size of the user stack.

**stripeswapflag**
(longword) striped swapping/paging area: 0 – disable; 1 – enable. Default is 0. Use the administration command /usr/alliant/extendswap to change this variable.

**swdevt** (longword; information only) disk swapping file space in units of 4096 bytes.

**swdevt+0x4**
(longword; information only) available disk swapping file space in units of 4096 bytes.

**swpgratio** (longword) percent of swapping/paging area used for swapping. Default is 45.

**sys_prefetch**
(longword) enables/disables prefetching on the VMEbusses. Default is ON. Prefetching allows the VMEbus controllers to achieve much higher data rates when they are reading out of global memory. It has no effect on the rates writing into memory.

**tcp_recvspace**
(longword) changes the buffer space the kernel allocates to tcp. Default is 0x1000. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**tcp_sendspace**
(longword) changes the buffer space the kernel allocates to tcp. Default is 0x1000. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**tz+4** (longword) daylight savings time: 1 – will switch to; 0 – will not switch to. Default is 1. Use the administration command /usr/alliant/timezone to set this variable.

**udp_recevspace**
(longword) changes the buffer space the kernel allocates to udp. Default is 0x4670. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**udp_sendspace**
(longword) changes the buffer space the kernel allocates to udp. Default is 0x2328. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**unpst_recvspace**
(longword) changes the buffer space the kernel allocates to pipes. Default is 0x4000. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**unpst_sendspace**
(longword) changes the buffer space the kernel allocates to pipes. Default is 0x4000. In general, increasing the buffer space improves throughput at the expense of using more cluster mbufs (see nmbclusters). Do not make bigger than 0x8000.

**usermem** (longword; information only) total user memory in units of 4096 bytes. Do not change.

**usrstack** (longword; information only) upper bound of user stack. Use usrstacksegs to change the size of the user stack.

**usrstacksegs**
(longword) size of the user stack in segments; the first segment is 3,964 kilobytes and the remaining segments are exactly 4 megabytes each. The minimum is 1 segment (3,964 kilobytes) and the maximum is 25 segments (102,268 kilobytes). Defaults to 19 (75,964 kilobytes).

**versionname**
(string; information only) name of system. Default is Alliant Concentrix-2800 and the version number.

## COMMANDS

| | |
|---|---|
| adb* | assembler level debugger |

*See *Concentrix-2800 Systems Programming Guide*

## FILES

| | |
|---|---|
| /dev/kmem | kernel memory |
| /unix | system kernel |

# Recover_system

## recover from hang or panic

Two categories of problems can stop operations on your system:

- **Hang** – The system fails to respond to commands from the console and terminals.
- **Panic** – The system types the word panic followed by a message (the panic string) on the console and invokes ddt, the kernel debugger. To terminals other than the console, the system appears to be hanging.

### LOG PROBLEM AND CALL CUSTOMER SERVICE

Call customer service (see Call_cs) before changing the state of the machine, if possible, to afford Alliant the opportunity to examine the problem first-hand. If customer service cannot be reached (for example, due to the hour) and the machine must be restarted, call customer service as soon as possible. Log the problem and your follow-on actions.

### RESTART FROM HANG

If the system does not respond to the console, your only option is to reboot Concentrix-2800; see Procedure E of Boot_system. If you cannot connect to Diagnostix-2800, reset the machine from the front panel; see Procedure C of Boot_system. If the reset procedure fails, power the machine off and on; see Procedure D of Boot_system.

### Page/Swap File Space Problems

The following options are now available in DDT for handling this type of system hangs without rebooting. These options allow you to select and kill specific processes in order to regain space. A brief guide for their use is included here. See the chapter on DDT in *Concentrix-2800 Systems Programming Guide* for complete information on using DDT.

- **:c kern_ps**
- **:d ps**
- **:c kill_proc**

If your system is hung or thrashing and you suspect that the load on some memory related resource is too high:

1. Hit the break key on the console to get into ddt.
2. From ddt, perform :c kern_ps or :d ps to select killable processes. See examples below.

```
------------------------------------------------------------------------------
:c kern_ps
                 TTY
  PID    UID  Maj/Minor  MM/S  SZPT      RSS  CURPGSPC  NSWP   NFAULTS  ARGS
    0      0  ( 0,   0)  1/1     0        0         0     0         0
    1      0  ( 0,   0)  2/1*    4        9         0     2         0  init @0
.
. removed for brevity
.
   65      0  ( 0,   0)  2/3*    4       18         0     2         0  /etc/syslogd
More?
                 TTY
  PID    UID  Maj/Minor  MM/S  SZPT      RSS  CURPGSPC  NSWP   NFAULTS  ARGS
   70      0  (c1,   0)  0/1*    4       20         0     4         0  /usr/lib/batch/b
   71      0  (c1,   0)  2/1*    4       15         0     1         0  ./netdae -n 2 /u
   77      0  ( 0,   0)  2/1*    4        7         0     1         0  /etc/nfsd 4
.
. removed for brevity
.
  109      0  (c1,   0)  18/1   12     2308      3634     0        57  /usr/cs/memhog 3
  110      0  (c1,   0)  8/3*   12     2330      3863     0        91  /usr/cs/memhog 3
  111      0  (c1,   0)  8/1    12     2728      3654     0       121  /usr/cs/memhog 3
40 processes scanned, 19 seem eligible for 2872 pages
[ip00] _kern_ps returned 0x13   $p
------------------------------------------------------------------------------
fields:
        PID             process pid
        UID             process uid
        TTY             maj/minor number of controlling tty (hex)
        MM              memory management flags (hex)
        S               process state (hex)
        SZPT            size of page table
        RSS             resident set size
        CURPGSPC        paging file consumption
        NSWP            number of swaps on this process
        NFAULTS         number of page faults incurred by this process;
                        all numbers      decimal, unless otherwise noted.
```

kern_ps will mark processes which are currently (read: right now) killable via a signal which will release some physical memory. The mark is indicated by an asterick after the process state field (for example process 110 above). Killability is determined by whether a process is currently:

- not sleeping or waiting at negative priority (i.e. won't receive the signal from kill_proc())
- consuming physical memory
- is not an idle job

Note that killability means that the process is immediately ready to accept a signal. A process not marked as immediately killable may be able to receive a signal in the near future.

alternately, directly from ddt:

```
------------------------------------------------------------------------------
:d ps
 UID   PID  PPID  C/P/NP SLEEP                COMMAND
   0     0     0  4/3/1  _sched+0x3C          Swapper
   0     1     0  3/0/1  _wait1+0x424         init D4
   0     2     0  4/0/1  _pagereclaimer+0x4C  Page reclaimer
   0     3     1  3/0/1  _ttread+0x1B8        -
   0     4     0  4/0/1  _cleaner+0x250       Primary page cleaner
.
. removed for brevity
.
   0    11     0 33/8/1  0x0                 ( Idle )
   0    12     0 33/8/1  0x0                 ( Idle )
   0    13     0 33/8/1  0x0                 ( Idle )
   0    14     0 33/8/1  ** CL0 **           ( Idle )
   0    15     0 33/8/1  ** CE2 **           ( Idle )
------------------------------------------
```

3. Select processes to kill: Select processes which are consuming large amounts of physical memory and/or paging file space. In the :c kern_ps example, these are processes 109, 110 and 111.

```
:c kill_proc 109.
[ip00] _kill_proc returned 0x0
:c kill_proc 110.
[ip00] _kill_proc returned 0x0
:c kill_proc 111.
[ip00] _kill_proc returned 0x0
```

Note the decimal point. It is very important.

### Methods of Avoiding Page/Swap File Space Problems

If the paging file space is running low, the console prints messages warning about the situation. To avoid a hang for this reason, take the following steps:

1. Increase the size of the page/swap file. The rule of thumb is as follows:

```
(size of job) < (physical memory size)
                + (paging portion of page file size)
                - 60 MB (for kernel, page tables, etc.)
```

You might also include enough additional space in the page/swap file for the other jobs you expect to be running in the system.

2. For cases where you know the need for paging space constantly arises, you can change the kernel variable *swpgratio* -- this variable is the percentage of page file space dedicated to swapping.

```
adb -w /unix /dev/kmem
swpgratio?W 20.     <--- this sets it to 20% (default %45)
CTRL d
```

## RESTART FROM PANIC

Unless customer service instructs you otherwise, you should record certain information and dump core (memory) before restarting the machine. A core dump is always performed on the SCSI Exabyte cartridge drive, which is always attached to machine 28 (IP0).

1. Turn on the console printer.

2. Type the following sequence of commands to ddt. The notation on the right explains the commands and is not part of the syntax. (The ESCAPE character echoes as a dollar sign.)

```
panicstr;RETURN          displays address of panic string
ESCAPE"\                 displays panic string
:pRETURN                 displays pending messages
:mRETURN                 displays machine states
:sRETURN                 displays stack of panic machine
:rRETURN                 displays registers of panic machine
:c dproc u@RETURN        displays proc data
:c du 4@RETURN           displays u data
```

3. Load a cartridge.

   Dump memory by typing the following:

```
28.ESCAPEESCAPEm         select machine 28. (decimal point needed)
:c dumpcoreRETURN        dump memory; note the lowercase c
```

   A successful dump exits with a status of 0. A status of 0x1 indicates an error (tape off-line, tape write-protected, data parity error). Fix the error, if possible, and retry the dump.

4. If the panic string contains the word "proceedable" or ends with a question mark, reactivate the system by typing ESCAPEp. If you get a prompt, initiate an orderly reboot (/etc/shutdown -r).

   Otherwise, reboot the system by connecting to FX/2800 startup mode (type CTRL^. on the console) and typing yRETURN (or selecting another reboot option; type ?RETURN to get a list of options).

# Remove_user

## remove a user from the system

Update the appropriate files so that the system no longer recognizes a user. Remove the files owned by that user.

### PROCEDURE

Log in as root and do the following:

1. Disable the user account by changing the password (passwd command):

```
# passwd hacker
New password: <password not echoed>
Retype new password: <password not echoed>
```

2. Save the user's files, as desired, by archiving them to tape (tar command) or copying them to another user's directory. If you copy the files to another user's directory, change the owner and group (if necessary) attributes of the file. Remove the user's files and login directory. The following example copies some of the files belonging to hacker to a subdirectory of gorn, then deletes all of the files in the directory tree starting at the login directory of hacker:

```
# cd ~hacker
# cp xyz_project/* ~gorn/
# chown gorn *
# cd ..
# rm -r hacker
```

3. Using a standard editor, remove the user's name from all entries on which it appears in /usr/lib/libaliases. If the user's name is the only name, remove the entry.

4. Type the following command:

```
/usr/ucb/newaliases
```

5. Remove the user's entry from /usr/spool/mail. You may want to first peruse the file (for example, with the more command, the mail -f command, or an editor).

The following example removes /usr/spool/mail/hacker:

```
% cd /usr/spool/mail
% more hacker
From gorn Sun Dec 13 19:35:50 1987
Received: by algoink.Algoink.COM (4.12/smail2.2/02-19-87)
        id AA04798; Sun, 13 Dec 87 19:35:42 est
Date: Sun, 13 Dec 87 19:35:42 est
From: gorn (Gorn)
Message-Id: <8712140035.AA04798@algoink.Algoink.COM>
To: hacker
Subject: Sorry to see you go
Status: RO
Too bad things didn't work out. Better luck on next job.
% rm hacker
```

6. Using a standard editor, remove the user's name from all entries on which it appears in /etc/group. If the user's name is the only name, remove the entry.

7. Locate any remaining files belonging to the user, archive them as desired, and delete them. The find command can be used to build a file in /tmp containing the names of the remaining files. The following example builds a file named /tmp/hacker containing the names of any files owned by hacker:

   ```
   # find / -user hacker -exec ls -lagd "{}" \; > /tmp/hacker&
   ```

   The command is usually executed in the background because it takes some time. When the job completes, examine the temporary file, dispose of any files named in the temporary file, and remove the temporary file.

   **Note:** On a system using the network file system (NFS), this command can significantly degrade interactive performance.

8. Using the command /etc/vipw, remove the entry in the file /etc/passwd for the user. By default, /etc/vipw places you in the vi editor; set the environment variable EDITOR to use a different editor (for example, setenv EDITOR emacs).

## NOTES

- This procedure must be modified if you are using network information service.

## COMMANDS

| | |
|---|---|
| cd* | change working directory |
| passwd -f* | change finger entry |
| chgrp* | change file group |
| chown | change owner |
| exit* | exit from shell |
| mkdir* | create directory |
| passwd* | change login password |
| newaliases | rebuild data base for mail aliases file |
| su* | substitute userid (superuser) |
| vipw | edit the password file |

*See *Concentrix-2800 Commands and Applications* or man1.

## FILES

| | |
|---|---|
| ~ | user login directory |
| /etc/group | group file |
| /etc/passwd | password file |
| /usr/lib/aliases | aliases file for sendmail |
| /usr/lib/aliases.dir | binary representation of aliases file |
| /usr/lib/aliases.pag | binary representation of aliases file |
| /usr/skel/.??* | skeleton files |

# Shutdown_system

## shut down system to single-user mode

Bring the system down to single-user mode from multi-user mode.

### PROCEDURE

Log in as root, and type the following:

**/etc/shutdown** *time* [*message*]RETURN

*Time* can be specified as now for immediate shutdown, *+n* for shutdown in *n* minutes, or *hour:minute* for shutdown at a certain hour (24-hour clock) and minute. *Message* can be anything and provides a means for warning users of the impending shutdown. Examples:

```
# /etc/shutdown +5 be back in five minutesRETURN

# /etc/shutdown 13:00 be back in five minutesRETURN

# /etc/shutdown nowRETURN
```

Specification of now on a system in use by multiple users is considered unfriendly.

From single-user mode, you can reboot the system or return directly to multi-user mode in the currently running system. See Procedure B of Boot_system.

You can also shut down by sending a terminate signal to the init process (see init in Appendix A), but this procedure should only be used in emergency and test situations.

### COMMANDS

| | |
|---|---|
| shutdown | shut down system at a given time |

### FILES

| | |
|---|---|
| /dunix | Diagnostix-2800 kernel image |
| /etc/nologin | file whose existence prevents logins |
| /fastboot | empty file whose existence signals a fastboot |
| /unix | Concentrix-2800 kernel image |
| /usr/adm/utmp | login record of current users |
| /usr/adm/wtmp | login record of all logins and logouts |
| /vmunix | a soft link to /unix |

# APPENDIX C
# SENDMAIL

**Note:** In general, the information in this appendix is not needed unless you are modifiying the sendmail configuration supplied by Alliant.

Sendmail implements a general internetwork mail routing facility, featuring aliasing and forwarding, automatic routing to network gateways, and flexible configuration.

In a simple network, each node has an address, and resources can be identified with a host-resource pair; in particular, the mail system can refer to users using a host-username pair. Host names and numbers have to be administered by a central authority, but usernames can be assigned locally to each host.

In an internet, multiple networks with different characterstics and managements must communicate. In particular, the syntax and semantics of resource identification change. Certain special cases can be handled trivially by *ad hoc* techniques, such as providing network names that appear local to hosts on other networks, as with the Ethernet at Xerox PARC. However, the general case is extremely complex. For example, some networks require point-to-point routing, which simplifies the database update problem since only adjacent hosts must be entered into the system tables, while others use end-to-end addressing. Some networks use a left-associative syntax and others use a right-associative syntax, causing ambiguity in mixed addresses.

Internet standards seek to eliminate these problems. Initially, these proposed expanding the address pairs to address triples, consisting of {network, host, resource} triples. Network numbers must be universally agreed upon, and hosts can be assigned locally on each network. The user-level presentation was quickly expanded to address domains, comprised of a local resource identification and a hierarchical domain specification with a common static root. The domain technique separates the issue of physical versus logical addressing. For example, an address of the form "eric@a.cc.berkeley.arpa" describes only the logical organization of the address space.

Sendmail is intended to help bridge the gap between the totally *ad hoc* world of networks that know nothing of each other and the clean, tightly-coupled world of unique network numbers. It can accept old arbitrary address syntaxes, resolving ambiguities using heuristics specified by the system administrator, as well as domain-based addressing. It helps guide the conversion of message formats between disparate networks. In short, sendmail is designed to assist a graceful transition to consistent internetwork addressing schemes.

## C.1 DESIGN GOALS

Design goals for sendmail include:

1. Compatibility with the existing mail programs, including Bell version 6 mail, Bell version 7 mail [UNIX83], Berkeley *Mail (System User's Guide)*, BerkNet mail[Schmidt79], and hopefully UUCP mail (Chapters 10 and 11). ARPANET mail [Crocker77a, Postel77] was also required.

2. Reliability, in the sense of guaranteeing that every message is correctly delivered or at least brought to the attention of a human for correct disposal; no message should ever be completely lost. This goal was considered essential because of the emphasis on mail in our environment. It has turned out to be one of the hardest goals to satisfy, especially in the face of the many anomalous message formats produced by various ARPANET sites. For example, certain sites generate improperly formated addresses, occasionally causing error-message loops. Some hosts use blanks in names, causing problems with UNIX mail programs that assume that an address is one word. The semantics of some fields are interpreted slightly differently by different sites. In summary, the obscure features of the ARPANET mail protocol really *are* used and are difficult to support, but must be supported.

3. Existing software to do actual delivery should be used whenever possible. This goal derives as much from political and practical considerations as technical.

4. Easy expansion to fairly complex environments, including multiple connections to a single network type (such as with multiple UUCP or Ether nets [Metcalfe76]). This goal requires consideration of the contents of an address as well as its syntax in order to determine which gateway to use. For example, the ARPANET is bringing up the TCP protocol to replace the old NCP protocol. No host at Berkeley runs both TCP and NCP, so it is necessary to look at the ARPANET host name to determine whether to route mail to an NCP gateway or a TCP gateway.

5. Configuration should not be compiled into the code. A single compiled program should be able to run as is at any site (barring such basic changes as the CPU type or the operating system). We have found this seemingly unimportant goal to be critical in real life. Besides the simple problems that occur when any program gets recompiled in a different environment, many sites like to "fiddle" with anything that they will be recompiling anyway.

6. Sendmail must be able to let various groups maintain their own mailing lists, and let individuals specify their own forwarding, without modifying the system alias file.

7. Each user should be able to specify which mailer to execute to process mail being delivered for him. This feature allows users who are using specialized mailers that use a different format to build their environment without changing the system, and facilitates specialized functions (such as returning an "I am on vacation" message).

8. Network traffic should be minimized by batching addresses to a single host where possible, without assistance from the user.

These goals motivated the architecture illustrated in Figure A-1. The user interacts with a mail generating and sending program. When the mail is created, the generator calls sendmail, which routes the message to the correct mailer(s). Since some of the senders can be network servers and some of the mailers can be network clients, sendmail can be used as an internet mail gateway.

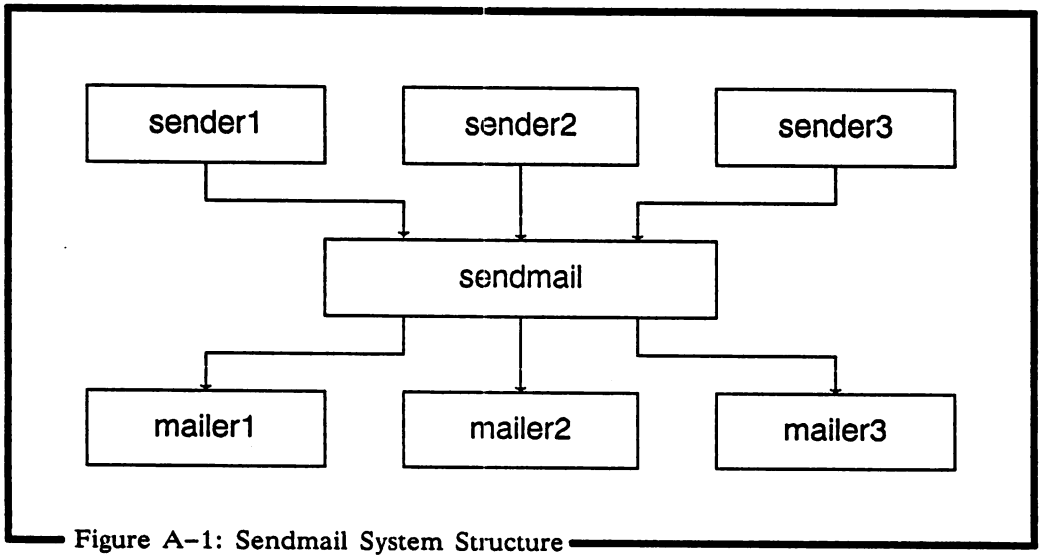


Figure A-1: Sendmail System Structure

## C.2 OVERVIEW

This section presents a general overview of the sendmail general internetwork mail routing facility.

### C.2.1 System Organization

Sendmail neither interfaces with the user nor does actual mail delivery. Rather, it collects a message generated by a user interface program (UIP) such as Berkeley *Mail*, MS [Crocker77b], or MH [Borden79], edits the message as required by the destination network, and calls appropriate mailers to do mail delivery or queueing for network transmission*. This discipline allows the insertion of new mailers at minimum cost. In this sense sendmail resembles the Message Processing Module (MPM) of [Postel79b].

* except when mailing to a file, when sendmail does the delivery directly.

## C.2.2 Interfaces to the Outside World

There are three ways sendmail can communicate with the outside world, both in receiving and in sending mail. These are using the conventional UNIX argument vector/return status, speaking SMTP over a pair of UNIX pipes, and speaking SMTP over an interprocess(or) channel.

### C.2.2.1 Argument Vector/exit Status

This technique is the standard UNIX method for communicating with the process. A list of recipients is sent in the argument vector, and the message body is sent on the standard input. Anything that the mailer prints is simply collected and sent back to the sender if there were any problems. The exit status from the mailer is collected after the message is sent, and a diagnostic is printed if appropriate.

### C.2.2.2 SMTP Over Pipes

The SMTP protocol [Postel82] can be used to run an interactive lock-step interface with the mailer. A subprocess is still created, but no recipient addresses are passed to the mailer via the argument list. Instead, they are passed one at a time in commands sent to the processes standard input. Anything appearing on the standard output must be a reply code in a special format.

### C.2.2.3 SMTP Over an IPC Connection

This technique is similar to the previous technique, except that it uses a 4.2BSD IPC channel [UNIX83]. This method is exceptionally flexible in that the mailer need not reside on the same machine. It is normally used to connect to a sendmail process on another machine.

## C.2.3 Operational Description

When a sender wants to send a message, it issues a request to *sendmail* using one of the three methods described above. Sendmail operates in two distinct phases. In the first phase, it collects and stores the message. In the second phase, message delivery occurs. If there were errors during processing during the second phase, sendmail creates and returns a new message describing the error and/or returns an status code telling what went wrong.

### C.2.3.1 Argument Processing and Address Parsing

If sendmail is called using one of the two subprocess techniques, the arguments are first scanned and option specifications are processed. Recipient addresses are then collected, either from the command line or from the SMTP RCPT command, and a list of recipients is created. Aliases are expanded at this step, including mailing lists. As much validation as possible of the addresses is done at this step: syntax is checked, and local addresses are verified, but detailed checking of host names and addresses is deferred until delivery. Forwarding is also performed as the local addresses are verified.

Sendmail appends each address to the recipient list after parsing. When a name is aliased or forwarded, the old name is retained in the list, and a flag is set that tells the delivery phase to ignore this recipient. This list is kept free from duplicates, preventing alias loops and duplicate messages deliverd to the same recipient, as might occur if a person is in two groups.

### C.2.3.2 Message Collection

Sendmail then collects the message. The message should have a header at the beginning. No formatting requirements are imposed on the message except that they must be lines of text (i.e., binary data is not allowed). The header is parsed and stored in memory, and the body of the message is saved in a temporary file.

To simplify the program interface, the message is collected even if no addresses were valid. The message is returned with an error.

### C.2.3.3 Message Delivery

For each unique mailer and host in the recipient list, sendmail calls the appropriate mailer. Each mailer invocation sends to all users receiving the message on one host. Mailers that only accept one recipient at a time are handled properly.

The message is sent to the mailer using one of the same three interfaces used to submit a message to sendmail. Each copy of the message is prepended by a customized header. The mailer status code is caught and checked, and a suitable error message given as appropriate. The exit code must conform to a system standard or a generic message ("Service unavailable") is given.

### C.2.3.4 Queueing for Retransmission

If the mailer returned an status that indicated that it might be able to handle the mail later, sendmail queues the mail and tries again later.

### C.2.3.5 Return to Sender

If errors occur during processing, sendmail returns the message to the sender for retransmission. The letter can be mailed back or written in the file dead.letter in the sender's home directory*.

---

* Obviously, if the site giving the error is not the originating site, the only reasonable option is to mail back to the sender. Also, there are many more error disposition options, but they only affect the error message – the "return to sender" function is always handled in one of these two ways.

### C.2.4 Message Header Editing

Certain editing of the message header occurs automatically. Header lines can be inserted under control of the configuration file. Some lines can be merged; for example, a "From:" line and a "Full-name:" line can be merged under certain circumstances.

### C.2.5 Configuration File

Almost all configuration information is read at runtime from an ASCII file, encoding macro definitions (defining the value of macros used internally), header declarations (telling sendmail the format of header lines that it will process specially, i.e., lines that it will add or reformat), mailer definitions (giving information such as the location and characteristics of each mailer), and address rewriting rules (a limited production system to rewrite addresses that is used to parse and rewrite the addresses).

To improve performance when reading the configuration file, a memory image can be provided. This provides a "compiled" form of the configuration file.

## C.3 USAGE AND IMPLEMENTATION

### C.3.1 Arguments

Arguments can be flags and addresses. Flags set various processing options. Following flag arguments, address arguments can be given, unless we are running in SMTP mode. Addresses follow the syntax in RFC822 [Crocker82] for ARPANET address formats. In brief, the format is:

1. Anything in parentheses is thrown away (as a comment).

2. Anything in angle brackets ("<>") is preferred over anything else. This rule implements the ARPANET standard that addresses of the form:

   `user name <machine-address>`

   sends to the electronic "machine-address" rather than the human "user name."

3. Double quotes ( " ) quote phrases; backslashes quote characters. Backslashes are more powerful in that they cause otherwise equivalent phrases to compare differently – for example, *user* and "*user*" are equivalent, but *user* is different from either of them.

Parentheses, angle brackets, and double quotes must be properly balanced and nested. The rewriting rules control remaining parsing*.

* Disclaimer: Some special processing is done after rewriting local names; see below.

## C.3.2 Mail to Files and Programs

Files and programs are legitimate message recipients. Files provide archival storage of messages, useful for project administration and history. Programs are useful as recipients in a variety of situations, for example, to maintain a public repository of systems messages (such as the Berkeley msgs program, or the MARS system [Sattley78]).

Any address passing through the initial parsing algorithm as a local address (i.e, not appearing to be a valid address for another mailer) is scanned for two special cases. If prefixed by a vertical bar ("|") the rest of the address is processed as a shell command. If the user name begins with a slash mark ("/") the name is used as a file name, instead of a login name.

Files that have setuid or setgid bits set but no execute bits set have those bits honored if sendmail is running as root.

## C.3.3 Aliasing, Forwarding, Inclusion

Sendmail reroutes mail three ways: Aliasing applies system wide. Forwarding allows each user to reroute incoming mail destined for that account. Inclusion directs sendmail to read a file for a list of addresses, and is normally used in conjunction with aliasing.

### C.3.3.1 Aliasing

Aliasing maps names to address lists using a system-wide file. This file is indexed to speed access. Only names that parse as local are allowed as aliases; this guarantees a unique key (since there are no nicknames for the local host).

### C.3.3.2 Forwarding

After alising, recipients that are local and valid are checked for the existence of a .forward file in their home directory. If it exists, the message is *not* sent to that user, but rather to the list of users in that file. Often this list contains only one address, and the feature is used for network mail forwarding.

Forwarding also permits a user to specify a private incoming mailer. For example, forwarding to:

```
"|/usr/local/newmail myname"
```

uses a different incoming mailer.

### C.3.3.3 Inclusion

Inclusion is specified in RFC 733 [Crocker77a] syntax:

```
:Include: pathname
```

An address of this form reads the file specified by *pathname* and sends to all users listed in that file.

The intent is *not* to support direct use of this feature, but rather to use this as a subset of aliasing. For example, an alias of the form:

```
project: :include:/usr/project/userlist
```

is a method of letting a project maintain a mailing list without interaction with the system administration, even if the alias file is protected.

It is not necessary to rebuild the index on the alias database when a :include: list is changed.

## C.3.4 Message Collection

Once all recipient addresses are parsed and verified, the message is collected. The message comes in two parts: a message header and a message body, separated by a blank line.

The header is formatted as a series of lines of the form

```
field-name: field-value
```

Field-value can be split across lines by starting the following lines with a space or a tab. Some header fields have special internal meaning, and have appropriate special processing. Other headers are simply passed through. Some header fields can be added automatically, such as time stamps.

The body is a series of text lines. It is completely uninterpreted and untouched, except that lines beginning with a dot have the dot doubled when transmitted over an SMTP channel. This extra dot is stripped by the receiver.

## C.3.5 Message Delivery

The send queue is ordered by receiving host before transmission to implement message batching. Each address is marked as it is sent so rescanning the list is safe. An argument list is built as the scan proceeds. Mail to files is detected during the scan of the send list. The interface to the mailer is performed using one of the techniques described in Section C.2.2.

After a connection is established, sendmail makes the per-mailer changes to the header and sends the result to the mailer. If any mail is rejected by the mailer, a flag is set to invoke the return-to-sender function after all delivery completes.

## C.3.6 Queued Messages

If the mailer returns a "temporary failure" exit status, the message is queued. A control file is used to describe the recipients to be sent to and various other parameters. This control file is formatted as a series of lines, each describing a sender, a recipient, the time of submission, or some other salient parameter of the message. The header of the message is stored in the control file, so that the associated data file in the queue is just the temporary file that was originally collected.

## C.3.7 Configuration

Configuration is controlled primarily by a configuration file read at startup. Sendmail should not need to be recomplied except

1. To change operating systems (V6, V7/32V, 4BSD).
2. To remove or insert the DBM (UNIX database) library.
3. To change ARPANET reply codes.
4. To add headers fields requiring special processing.

Adding mailers or changing parsing (i.e., rewriting) or routing information does not require recompilation.

If the mail is being sent by a local user, and the file .mailcf exists in the sender's home directory, that file is read as a configuration file after the system configuration file. The primary use of this feature is to add header lines.

The configuration file encodes macro definitions, header definitions, mailer definitions, rewriting rules, and options.

### C.3.7.1 Macros

Macros can be used in three ways. Certain macros transmit unstructured textual information into the mail system, such as the name sendmail uses to identify itself in error messages. Other macros transmit information from sendmail to the configuration file for use in creating other fields (such as argument vectors to mailers); e.g., the name of the sender, and the host and user of the recipient. Other macros are unused internally, and can be used as shorthand in the configuration file.

### C.3.7.2 Header Declarations

Header declarations inform sendmail of the format of known header lines. Knowledge of a few header lines is built into sendmail, such as the "From:" and "Date:" lines.

Most configured headers are automatically inserted in the outgoing message if they don't exist in the incoming message. Certain headers are suppressed by some mailers.

### C.3.7.3 Mailer Declarations

Mailer declarations tell sendmail of the various mailers available to it. The definition specifies the internal name of the mailer, the pathname of the program to call, some flags associated with the mailer, and an argument vector to be used on the call; this vector is macro-expanded before use.

### C.3.7.4 Address Rewriting Rules

The heart of address parsing in sendmail is a set of rewriting rules. These are an ordered list of pattern-replacement rules, (somewhat like a production system, except that order is critical), that are applied to each address. The address is rewritten textually until it is either rewritten into a special canonical form (i.e., a (mailer, host, user) 3-tuple, such as {arpanet, usc-isif, postel} representing the address "postel@uscisif"), or it falls off the end. When a pattern matches, the rule is reapplied until it fails.

The configuration file also supports the editing of addresses into different formats. For example, an address of the form:

```
ucsfcgl!tef
```

might be mapped into:

```
tef@ucsfcgl.UUCP
```

to conform to the domain syntax. Translations can also be done in the other direction.

### C.3.7.5 Option Setting

There are several options that can be set from the configuration file. These include the pathnames of various support files, timeouts, default modes, etc.

## C.4 COMPARISON WITH OTHER MAILERS

### C.4.1 Delivermail

Sendmail is an outgrowth of delivermail. The primary differences are:

1. Configuration information is not compiled in. This change simplifies many of the problems of moving to other machines. It also allows easy debugging of new mailers.

2. Address parsing is more flexible. For example, delivermail only supported one gateway to any network, whereas sendmail can be sensitive to host names and reroute to different gateways.

3. Forwarding and :include: features eliminate the requirement that the system alias file be writable by any user (or that an update program be written, or that the system administration make all changes).

4. Sendmail supports message batching across networks when a message is being sent to multiple recipients.

5. A mail queue is provided in sendmail. Mail that cannot be delivered immediately but can potentially be delivered later is stored in this queue for a later retry. The queue also provides a buffer against system crashes; after the message has been collected it can be reliably redelivered even if the system crashes during the initial delivery.

6. Sendmail uses the networking support provided by 4.2BSD to provide a direct interface networks such as the ARPANET and/or Ethernet using SMTP (the Simple Mail Transfer Protocol) over a TCP/IP connection.

## C.4.2 MMDF

MMDF [Crocker79] spans a wider problem set than sendmail. For example, the domain of MMDF includes a phone network mailer, whereas sendmail calls on preexisting mailers in most cases.

MMDF and sendmail both support aliasing, customized mailers, message batching, automatic forwarding to gateways, queueing, and retransmission. MMDF supports two-stage timeout, which sendmail does not support.

The configuration for MMDF is compiled into the code*.

Since MMDF does not consider backwards compatibility as a design goal, the address parsing is simpler but much less flexible.

It is somewhat harder to integrate a new channel** into MMDF. In particular, MMDF must know the location and format of host tables for all channels, and the channel must speak a special protocol. This allows MMDF to do additional verification (such as verifying host names) at submission time.

MMDF strictly separates the submission and delivery phases. Although sendmail has the concept of each of these stages, they are integrated into one program, whereas in MMDF they are split into two programs.

## C.4.3 Message Processing Module

The Message Processing Module (MPM) discussed by Postel [Postel79b] matches sendmail closely in terms of its basic architecture. However, like MMDF, the MPM includes the network interface software as part of its domain.

MPM also postulates a duplex channel to the receiver, as does MMDF, thus allowing simpler handling of errors by the mailer than is possible in sendmail. When a message queued by sendmail is sent, any errors must be returned to the sender by the mailer

* Dynamic configuration tables are currently being considered for MMDF; allowing the installer to select either compiled for dynamic tables.

** The MMDS equivalent of a sendmail "mailer."

itself. Both MPM and MMDF mailers can return an immediate error response, and a single error processor can create an appropriate response.

MPM prefers passing the message as a structured object, with type-length-value tuples*. Such a convention requires a much higher degree of cooperation between mailers than is required by sendmail. MPM also assumes a universally agreed upon internet name space (with each address in the form of a net-hostuser tuple), which sendmail does not.

## C.5 EVALUATIONS AND FUTURE PLANS

Sendmail is designed to work in a nonhomogeneous environment. Every attempt is made to avoid imposing unnecessary constraints on the underlying mailers. This goal has driven much of the design. One of the major problems has been the lack of a uniform address space, as postulated in [Postel79a] and [Postel79b].

A nonuniform address space implies that a path is specified in all addresses, either explicitly (as part of the address) or implicitly (as with implied forwarding to gateways). This restriction has the unpleasant effect of making replying to messages exceedingly difficult, since there is no one "address" for any person, but only a way to get there from wherever you are.

Interfacing to mail programs that were not initially intended to be applied in an internet environment has been amazingly successful, and has reduced the job to a manageable task.

Sendmail has knowledge of a few difficult environments built in. It generates ARPANET FTP/SMTP compatible error messages (prepended with three-digit numbers [Neigus73, Postel74, Postel82]) as necessary, optionally generates UNIX-style "From" lines on the front of messages for some mailers, and knows how to parse the same lines on input. Also, error handling has an option customized for BerkNet.

The decision to avoid doing any type of delivery where possible (even, or perhaps especially, local delivery) has turned out to be a good idea. Even with local delivery, there are issues of the location of the mailbox, the format of the mailbox, the locking protocol used, etc., that are best decided by other programs. One surprisingly major annoyance in many internet mailers is that the location and format of local mail is built in. The feeling seems to be that local mail is so common that it should be efficient. This feeling is not born out by our experience; on the contrary, the location and format of mailboxes seems to vary widely from system to system.

The ability to automatically generate a response to incoming mail (by forwarding mail to a program) seems useful ("I am on vacation until late August....") but can create problems such as forwarding loops (two people on vacation whose programs send notes back and forth, for instance) if these programs are not well written. A program could be written to do standard tasks correctly, but this would solve the general case.

It might be desirable to implement some form of load limiting. I am unaware of any mail system that addresses this problem, nor am I aware of any reasonable solution at this time.

* This is similar to the NBS standard.

The configuration file is currently practically inscrutable; considerable convenience could be realized with a higher-level format.

It seems clear that common protocols will be changing soon to accommodate changing requirements and environments. These changes will include modifications to the message header (e.g., [NBS80]) or to the body of the message itself (such as for multimedia messages [Postel80]). Experience indicates that these changes should be relatively trivial to integrate into the existing system.

In tightly coupled environments, it would be nice to have a name server such as Grapvine [Birrell82] integrated into the mail system. This would allow a site such as "Berkeley" to appear as a single host, rather than as a collection of hosts, and would allow people to move transparently among machines without having to change their addresses. Such a facility would require an automatically updated database and some method of resolving conflicts. Ideally this would be effective even without all hosts being under a single management. However, it is not clear whether this feature should be integrated into the aliasing facility or should be considered a "value added" feature outside sendmail itself.

As a more interesting case, the CSNET name server [Solomon81] provides an facility that goes beyond a single tightly-coupled environment. Such a facility would normally exist outside of sendmail however.

## C.6 REFERENCES


[Birrell82] Birrell, A. D., Levin, R., Needham, R. M., and Schroeder, M. D., "Grapevine: An Exercise in Distributed Computing." In *Comm. A.C.M. 25*, 4, April 82.

[Borden79] Borden, S., Gaines, R. S., and Shapiro, N. Z., *The MH Message Handling System: Users' Manual*. R-2367-PAF. Rand Corporation. October 1979.

[Crocker77a] Crocker, D. H., Vittal, J. J., Pogran, K. T., and Henderson, D. A. Jr., *Standard for the Format of ARPA Network Text Messages*. RFC 733, NIC 41952. In [Feinler78]. November 1977.

[Crocker77b] Crocker, D. H., *Framework and Functions of the MS Personal Message System*. R-2134-ARPA, Rand Corporation, Santa Monica, California. 1977.

[Crocker79] Crocker, D. H., Szurkowski, E. S., and Farber, D. J., *An Internetwork Memo Distribution Facility – MMDF. 6th Data Communication Symposium, Asilomar. November 1979.*

[Crocker82] Crocker, D. H., *Standard for the Format of Arpa Internet Text Messages*. RFC 822. Network Information Center, SRI International, Menlo Park, California. August 1982.

[Metcalfe76] Metcalfe, R., and Boggs, D., "Ethernet: Distributed Packet Switching for Local Computer Networks", *Communications of the ACM 19*, 7. July 1976.

[Feinler78] Feinler, E., and Postel, J. (eds.), *ARPANET Protocol Handbook*. NIC 7104, Network Information Center, SRI International, Menlo Park, California. 1978.

[NBS80] National Bureau of Standards, *Specification of a Draft Message Format Standard*. Report No. ICST/CBOS 80-2. October 1980.

[Neigus73] Neigus, N., *File Transfer Protocol for the ARPA Network*. RFC 542, NIC 17759. In [Feinler78]. August, 1973.

[Postel74] Postel, J., and Neigus, N., Revised FTP Reply Codes. RFC 640, NIC 30843. In [Feinler78]. June, 1974.

[Postel77] Postel, J., *Mail Protocol*. NIC 29588. In [Feinler78]. November 1977.

[Postel79a] Postel, J., *Internet Message Protocol*. RFC 753, IEN 85. Network Information Center, SRI International, Menlo Park, California. March 1979.

[Postel79b] Postel, J. B., *An Internetwork Message Structure. In Proceedings of the Sixth Data Communications Symposium, IEEE. New York. November 1979.*

[Postel80] Postel, J. B., *A Structured Format for Transmission of Multi-Media Documents*. RFC 767. Network Information Center, SRI International, Menlo Park, California. August 1980.

[Postel82] Postel, J. B., *Simple Mail Transfer Protocol*. RFC821 (obsoleting RFC788). Network Information Center, SRI International, Menlo Park, California. August 1982.

[Schmidt79] Schmidt, E., *An Introduction to the Berkeley Network*. University of California, Berkeley California. 1979.

[Sluizer81] Sluizer, S., and Postel, J. B., *Mail Transfer Protocol*. RFC 780. Network Information Center, SRI International, Menlo Park, California. May 1981.

[Solomon81] Solomon, M., Landweber, L., and Neuhengen, D., "The Design of the CSNET Name Server." CS-DN-2, University of Wisconsin, Madison. November 1981.

[Su82] Su, Zaw-Sing, and Postel, Jon, *The Domain Naming Convention for Internet User Applications. RFC819. Network Information Center, SRI International, Menlo Park, California. August 1982.*

[UNIX83] *The UNIX Programmer's Manual, Seventh Edition, Virtual VAX-11 Version, Volume 1. Bell Laboratories, modified by the University of California, Berkeley, California. March, 1983.*

## C.7 SENDMAIL INSTALLATION AND OPERATION

Sendmail implements a general purpose internetwork mail routing facility. It is not tied to any one transport protocol – its function can be likened to a crossbar switch, relaying messages from one domain into another. In the process, it can do a limited amount of message header editing to put the message into a format that is appropriate for the receiving domain. All of this is done under the control of a configuration file.

Due to the requirements of flexibility for sendmail, the configuration file can seem somewhat unapproachable. However, there are only a few basic configurations for most sites, for which standard configuration files have been supplied. Most other configurations can be built by adjusting an existing configuration files incrementally.

Although sendmail is intended to run without the need for monitoring, it has a number of features that can be used to monitor or adjust the operation under unusual circumstances. The following features are described:

- Basic sendmail installation.
- The day-to-day information you should know to maintain your mail system. If you have a relatively normal site, this should contain sufficient information for you to install sendmail and keep it happy.
- Some parameters that can be safely tweaked.
- Command line arguments.
- Detailed information about the configuration file for people who must write their own.
- Brief but detailed explanations of a number of features not described in the rest of the chapter.

The references in this chapter can be found in Appendix NO TAG: *Sendmail – An Internetwork Mail Router*.

### C.7.1 Basic Installation

There are two basic steps to installing sendmail. The hard part is to build the configuration table. This is a file that sendmail reads when it starts up that describes the mailers it knows about, how to parse addresses, how to rewrite the message header, and the settings of various options. Although the configuration table is quite complex, a configuration can usually be built by adjusting an existing off-the-shelf configuration. The second part is actually doing the installation, i.e., creating the necessary files, etc.

The remainder of this section describes the installation of sendmail assuming you can use one of the existing configurations and that the standard installation parameters are acceptable. All pathnames and examples are given from the root of the sendmail subtree.

### C.7.1.1 Off-the-shelf Configurations

The configuration files are all in the subdirectory cf of the sendmail directory. The ones used at Berkeley are in m4 format; files with names ending .m4 are m4 include files, while files with names ending .mc are the master files. Files with names ending .cf are the m4 processed versions of the corresponding .mc file.

Two off-the-shelf configuration files are supplied to handle the basic cases: cf/arpaproto.cf for Arpanet (TCP) sites and cf/uucpproto.cf for UUCP sites. These are *not* in m4 format. Copy the file you need to a file with the same name as your system. For example:

```
cp uucpproto.cf ucsfcgl.cf
```

This file is now ready for installation as /usr/lib/sendmail.cf.

### C.7.1.2 Installation Using the Makefile

A makefile exists in the root of the sendmail directory that does all of these steps for a 4.2BSD system. It may have to be slightly tailored for use on other systems.

Before using this makefile, you should already have created your configuration file and left it in the file cf/*system*.cf where *system* is the name of your system (i.e., what is returned by the hostname command). If you do not have *hostname* you can use the declaration "HOST=*system*" on the make command line. Also examine the file md/config.m4 and change the m4 macros there to reflect any libraries and compilation flags you may need.

The basic installation procedure is to type:

```
make
make install
```

in the root directory of the sendmail distribution. This makes all binaries and install them in the standard places. The second make command must be executed as the superuser (root).

### C.7.1.3 Installation by Hand

Along with building a configuration file, you must install the sendmail startup into your UNIX system. If you are doing this installation in conjunction with a regular Berkeley UNIX install, these steps are already complete. Many of these steps will have to be executed as the superuser (root).

#### lib/libsys.a

The library in lib/libsys.a contains some routines that should in some sense be part of the system library. These are the system logging routines and the new directory access routines (if required). If you are not running the new 4.2BSD directory code and do not have the compatibility routines installed in your system library, execute the commands:

```
cd lib
make ndir
```

This compiles and installs the 4.2 compatibility routines in the library. Then type:

```
cd lib# if required
make
```

This recompiles and fills the library.

#### /usr/lib/sendmail

The binary for sendmail is located in /usr/lib. There is a version available in the source directory that is probably inadequate for your system. Plan on recompiling and installing the entire system:

```
cd src
rm -f *.o
make
cp sendmail /usr/lib
```

#### /usr/lib/sendmail.cf

Install the configuration file that you created earlier in /usr/lib/sendmail.cf:

```
cp cf/system.cf /usr/lib/sendmail.cf
```

#### /usr/ucb/newaliases

If you are running delivermail, it is critical that the newaliases command be replaced. This can just be a link to sendmail:

```
rm -f /usr/ucb/newaliases
ln /usr/lib/sendmail /usr/ucb/newaliases
```

### /usr/lib/sendmail.cf

The configuration file must be installed in /usr/lib. This is described above.

### /usr/spool/mqueue

Create the directory /usr/spool/mqueue to hold the mail queue. This directory should be mode 777 unless sendmail is run setuid, when mqueue should be owned by the sendmail owner and mode 755.

### /usr/lib/aliases*

The system aliases are held in three files. The file /usr/lib/aliases is the master copy. A sample is given in lib/aliases which includes some aliases that *must* be defined:

```
cp lib/aliases /usr/lib/aliases
```

Extend this file with any aliases that are apropos to your system.

Normally sendmail looks at a version of these files maintained by the dbm routines. These are stored in /usr/lib/aliases.dir and /usr/lib/aliases.pag. These can initially be created as empty files, but they will have to be initialized promptly. These should be mode 666 if you are running a reasonably relaxed system:

```
cp /dev/null /usr/lib/aliases.dir
cp /dev/null /usr/lib/aliases.pag
chmod 666 /usr/lib/aliases.*
newaliases
```

### /etc/rc

It is necessary to start up the sendmail daemon when your system reboots. This daemon performs two functions: it listens on the SMTP socket for connections (to receive mail from a remote system) and it processes the queue periodically to insure that mail gets delivered when hosts come up.

Add the following lines to /etc/rc (or /etc/rc.local as appropriate) in the area where it is starting up the daemons:

```
if [ -f /usr/lib/sendmail ]; then
        (cd /usr/spool/mqueue; rm -f [lnx]f*)
        /usr/lib/sendmail -bd -q30m &
        echo -n ' sendmail' >/dev/console
fi
```

The cd and rm commands insure that all lock files have been removed; extraneous lock files may be left around if the system goes down in the middle of processing a message. The line that actually invokes sendmail has two flags: -bd causes it to listen on the SMTP port, and -q30m causes it to run the queue every half hour.

If you are not running a version of UNIX that supports Berkeley TCP/IP, do not include the -bd flag.

### /usr/lib/sendmail.hf

This is the help file used by the SMTP HELP command. Copy it from lib/sendmail.hf:

```
cp lib/sendmail.hf /usr/lib
```

### /usr/lib/sendmail.st

If you wish to collect statistics about your mail traffic, create the file /usr/lib/sendmail.st:

```
cp /dev/null /usr/lib/sendmail.st
chmod 666 /usr/lib/sendmail.st
```

This file does not grow. It is printed with the program aux/mailstats.

### /etc/syslog

You may want to run the syslog program (to collect log information about sendmail). This program normally resides in /etc/syslog, with support files /etc/syslog.conf and /etc/syslog.pid. The program is located in the aux subdirectory of the sendmail distribution. The file /etc/syslog.conf describes the file(s) that sendmail logs in. For a complete description of syslog, see the command syslog in this book.

### /usr/ucb/newaliases

If sendmail is invoked as "newaliases," it simulates the -bi flag (i.e., rebuilds the alias database; see below). This should be a link to /usr/lib/sendmail.

#### /usr/ucb/mailq

If sendmail is invoked as mailq, it simulates the -bp flag (i.e., sendmail prints the contents of the mail queue; see below). This should be a link to /usr/lib/sendmail.

## C.7.2 Normal Operations

### C.7.2.1 The System Log

The system log is supported by the syslog program.

#### Format

Each line in the system log consists of a timestamp, the name of the machine that generated it (for logging from several machines over the ethernet), the word "sendmail:", and a message.

#### Levels

If you have syslog or an equivalent installed, you can do logging. There is a large amount of information that can be logged. The log is arranged as a succession of levels. At the lowest level only extremely strange situations are logged. At the highest level, even the most mundane and uninteresting events are recorded for posterity. As a convention, log levels under ten are considered useful; log levels above ten are usually for debugging purposes.

A complete description of the log levels is given in Section C.7.4.3.

### C.7.2.2 The Mail Queue

The mail queue should be processed transparently. However, you may find that manual intervention is sometimes necessary. For example, if a major host is down for a period of time the queue may become clogged. Although sendmail ought to recover gracefully when the host comes up, you may find performance unacceptably bad in the meantime.

#### Printing the Queue

The contents of the queue can be printed using the mailq command (or by specifying the -bp flag to sendmail):

```
mailq
```

This produces a listing of the queue id's, the size of the message, the date the message entered the queue, and the sender and recipients.

### Format of Queue Files

All queue files have the form *xfAA99999* where *AA99999* is the *id* for this file and the *x* is a type. The types are:

d The data file. The message body (excluding the header) is kept in this file.

l The lock file. If this file exists, the job is currently being processed, and a queue run does not process the file. For that reason, an extraneous lf file can cause a job to apparently disappear (it does not even time out!).

n This file is created when an id is being created. It is a separate file to ensure that no mail can ever be destroyed due to a race condition. It should exist for no more than a few milliseconds at any given time.

q The queue control file. This file contains the information necessary to process the job.

t A temporary file. These are an image of the qf file when it is being rebuilt. It should be renamed to an qf file very quickly.

x A transcript file, existing during the life of a session showing everything that happens during that session.

The qf file is structured as a series of lines each beginning with a code letter. The lines are as follows:

D The name of the data file. There can only be one of these lines.

H A header definition. There can be any number of these lines. The order is important: they represent the order in the final message. These use the same syntax as header definitions in the configuration file.

R A recipient address. This is normally completely aliased, but is actually realiased when the job is processed. There is one line for each recipient.

S The sender address. There can only be one of these lines.

T The job creation time. This is used to compute when to time out the job.

P The current message priority. This is used to order the queue. Higher numbers mean lower priorities. The priority increases as the message sits in the queue. The initial priority depends on the message class and the size of the message.

M A message. This line is printed by the mailq command, and is generally used to store status information. It can contain any text.

As an example, the following is a queue file sent to mckusick@calder and wnj:

```
DdfA13557
Seric
T404261372 P132
Rmckusick@calder
Rwnj
H?D?date: 23-Oct-82 15:49:32-PDT (Sat)
H?F?from: eric (Eric Allman)
H?x?full-name: Eric Allman
Hsubject: this is an example message
Hmessage-id: <8209232249.13557@UCBARPA.BERKELEY.ARPA>
Hreceived: by UCBARPA.BERKELEY.ARPA (3.227 [10/22/82])
      id A13557; 23-Oct-82 15:49:32-PDT (Sat)
Hphone: (415) 548-3211
HTo: mckusick@calder, wnj
```

This shows the name of the data file, the person who sent the message, the submission time (in seconds since January 1, 1970), the message priority, the message class, the recipients, and the headers for the message.

### Forcing the Queue

Sendmail should run the queue automatically at intervals. The algorithm is to read and sort the queue, and then to attempt to process all jobs in order. When it attempts to run the job, sendmail first checks to see if the job is locked. If so, it ignores the job.

There is no attempt to insure that only one queue processor exists at any time, since there is no guarantee that a job cannot take forever to process. Due to the locking algorithm, it is impossible for one job to freeze the queue. However, an uncooperative recipient host or a program recipient that never returns can accumulate many processes in your system. Unfortunately, there is no way to resolve this without violating the protocol.

In some cases, you may find that a major host going down for a couple of days can create a prohibitively large queue. This results in sendmail spending an inordinate amount of time sorting the queue. This situation can be fixed by moving the queue to a temporary place and creating a new queue. The old queue can be run later when the offending host returns to service.

To do this, it is acceptable to move the entire queue directory:

```
cd /usr/spool
mv mqueue omqueue; mkdir mqueue; chmod 777 mqueue
```

Then kill the existing daemon (since it is processing in the old queue directory) and create a new daemon.

To run the old mail queue, run the following command:

```
/usr/lib/sendmail -oQ/usr/spool/omqueue -q
```

The -oQ flag specifies an alternate queue directory and the -q flag says to just run every job in the queue. You can use the -v flag to watch what is going on.

When the queue is finally emptied, you can remove the directory:

```
rmdir /usr/spool/omqueue
```

### C.7.2.3 The Alias Database

The alias database exists in two forms. One is a text form, maintained in the file /usr/lib/aliases. The aliases are of the form

```
name: name1, name2, ...
```

Only local names can be aliased; e.g.,

```
eric@mit-xx: eric@berkeley
```

does not have the desired effect. Aliases can be continued by starting any continuation lines with a space or a tab. Blank lines and lines beginning with a sharp sign ("#") are comments.

The second form is processed by the dbm library. This form is in the files /usr/lib/aliases.dir and /usr/lib/aliases.pag. This is the form that sendmail actually uses to resolve aliases. This technique is used to improve performance.

#### Rebuilding the Alias Database

The DBM version of the database can be rebuilt explicitly by executing the command

```
newaliases
```

This is equivalent to giving sendmail the -bi flag:

```
/usr/lib/sendmail -bi
```

If the D option is specified in the configuration, sendmail rebuilds the alias database automatically if possible when it is out of date. The conditions under which it does this are:

1. The DBM version of the database is mode 666.
2. Sendmail is running setuid to root.

Auto-rebuild can be dangerous on heavily loaded machines with large alias files; if it might take more than five minutes to rebuild the database, there is a chance that several processes will start the rebuild process simultaneously.

#### Potential Problems

There are a number of problems that can occur with the alias database. They all result from a sendmail process accessing the DBM version while it is only partially built. This can happen under two circumstances: One process accesses the database while another process is rebuilding it, or the process rebuilding the database dies (due to being killed or a system crash) before completing the rebuild.

Sendmail has two techniques to try to relieve these problems. First, it ignores interrupts while rebuilding the database; this avoids the problem of someone aborting the process leaving a partially rebuilt database. Second, at the end of the rebuild it adds an alias of the form

```
@: @
```

(which is not normally valid). Before sendmail accesses the database, it checks to insure that this entry exists*. It waits up to five minutes for this entry to appear, at which point it forces a rebuild itself**.

### List Owners

If an error occurs on sending to a certain address, say "*x*", sendmail looks for an alias of the form "owner-*x*" to receive the errors. This is typically useful for a mailing list where the submitter of the list has no control over the maintanence of the list itself; in this case the list maintainer would be the owner of the list. For example:

```
unix-wizards: eric@ucbarpa, wnj@monet, nosuchuser,
        sam@matisse
owner-unix-wizards: eric@ucbarpa
```

would cause "eric@ucbarpa" to get the error that occurs when someone sends to unix-wizards due to the inclusion of "nosuchuser" on the list.

## C.7.2.4 Per-User Forwarding (.forward Files)

As an alternative to the alias database, any user can put a file with the name .forward in his or her home directory. If this file exists, sendmail redirects mail for that user to the list of addresses listed in the .forward file. For example, if the home directory for user "mckusick" has a .forward file with contents:

```
mckusick@ernie
kirk@calder
```

any mail arriving for "mckusick" is redirected to the specified accounts.

---

* The "a" option is required in the configuration for this action to occur. This should normally be specified unless you are running delivermail in parallel with sendmail.

** The "D" option must be specified in the configuration file for this operation to occur.

### C.7.2.5 Special Header Lines

Several header lines have special interpretations defined by the configuration file. Others have interpretations built into sendmail that cannot be changed without changing the code. These builtins are described here.

#### Return-Receipt-To:

If this header is sent, a message is sent to any specified addresses when the final delivery is complete, if the mailer has the l flag (local delivery) set in the mailer descriptor.

#### Errors-To:

If errors occur anywhere during processing, this header causes error messages to go to the listed addresses rather than to the sender. This is intended for mailing lists.

#### Apparently-To:

If a message comes in with no recipients listed in the message (in a To:, Cc:, or Bcc: line), sendmail adds an "Apparently-To:" header line for any recipients it is aware of. This is not put in as a standard recipient line to warn any recipients that the list is not complete.

At least one recipient line is required under RFC 822.

## C.7.3 Arguments

The complete list of arguments to sendmail is described in detail in Section C.7.6. Some important arguments are described here.

### C.7.3.1 Queue Interval

The amount of time between forking a process to run through the queue is defined by the -q flag. If you run in mode f or a this can be relatively large, since it is only relevant when a host that was down comes back up. If you run in q mode it should be relatively short, since it defines the maximum amount of time that a message can sit in the queue.

### C.7.3.2 Daemon Mode

If you allow incoming mail over an IPC connection, you should have a daemon running. This should be set by your /etc/rc file using the -bd flag. The -bd flag and the -q flag can be combined in one call:

```
/usr/lib/sendmail -bd -q30m
```

### C.7.3.3 Forcing the Queue

In some cases you may find that the queue is clogged for some reason. You can force a queue run using the -q flag (with no value). It is entertaining to use the -v flag (verbose) when this is done to watch what happens:

```
/usr/lib/sendmail -q -v
```

### C.7.3.4 Debugging

There are a fairly large number of debug flags built into sendmail. Each debug flag has a number and a level, where higher levels means to print out more information. The convention is that levels greater than nine are absurd; they print out so much information that you wouldn't normally want to see them except for debugging that particular piece of code. Debug flags are set using the -d option; the syntax is:

```
debug-flag:  -d debug-list
debug-list:  debug-option [ , debug-option ]
debug-option:       debug-range [ . debug-level ]
debug-range: integer | integer - integer
debug-level: integer
```

where spaces are for reading ease only. For example,

```
-d12   Set flag 12 to level 1
-d12.3       Set flag 12 to level 3
-d3-17       Set flags 3 through 17 to level 1
-d3-17.4     Set flags 3 through 17 to level 4
```

For a complete list of the available debug flags, look at the code.

### C.7.3.5 Trying a Different Configuration File

An alternative configuration file can be specified using the -C flag; for example,

```
/usr/lib/sendmail -Ctest.cf
```

uses the configuration file test.cf instead of the default /usr/lib/sendmail.cf. If the -C flag has no value it defaults to sendmail.cf in the current directory.

### C.7.3.6 Changing the Values of Options

Options can be overridden using the -o flag. For example,

```
/usr/lib/sendmail -oT2m
```

sets the T (timeout) option to two minutes for this run only.

## C.7.4 Tuning

You may want to change a number of configuration parameters, depending on the requirements of your site. Most of these are set using an option in the configuration file. For example, the line "OT3d" sets option T to the value "3d" (three days).

### C.7.4.1 Timeouts

All time intervals are set using a scaled syntax. For example, "10m" represents ten minutes, whereas "2h30m" represents two and a half hours. The full set of scales is:

| | |
|---|---|
| s | seconds |
| m | minutes |
| h | hours |
| d | days |
| w | weeks |

#### Queue Interval

The argument to the -q flag specifies how often a subdaemon runs the queue. This is typically set to between five minutes and one half hour.

#### Read Timeouts

It is possible to time out when reading the standard input or when reading from a remote SMTP server. Technically, this is not acceptable within the published protocols. However, it might be appropriate to set it to something large in certain environments (such as an hour). This reduces the chance of large numbers of idle daemons piling up on your system. This timeout is set using the r option in the configuration file.

#### Message Timeouts

After sitting in the queue for a few days, a message times out. This is to insure that at least the sender is aware of the inability to send a message. The timeout is typically set to three days. This timeout is set using the T option in the configuration file.

The time of submission is set in the queue, rather than the amount of time left until timeout. As a result, you can flush messages that have been hanging for a short period by running the queue with a short message timeout. For example,

```
/usr/lib/sendmail -oT1d -q
```

runs the queue and flushes anything that is one day old.

### C.7.4.2 Delivery Mode

sendmail can operate in a number of delivery modes. These are set by the d configuration option. These modes specify how quickly mail is delivered. Valid modes are:

| | |
|---|---|
| i | deliver interactively (synchronously) |
| b | deliver in background (asynchronously) |
| q | queue only (don't deliver) |

There are tradeoffs. Mode i passes the maximum amount of information to the sender, but is hardly ever necessary. Mode q puts the minimum load on your machine, but means that delivery may be delayed for up to the queue interval. Mode b is probably a good compromise. However, this mode can cause large numbers of processes if you have a mailer that takes a long time to deliver a message.

### C.7.4.3 Log Level

The level of logging can be set for sendmail. The default using a standard configuration table is level 9. The levels are as follows:

| | |
|---|---|
| 0 | No logging. |
| 1 | Major problems only. |
| 2 | Message collections and failed deliveries. |
| 3 | Successful deliveries. |
| 4 | Messages being defered (due to a host being down, etc.). |
| 5 | Normal message queueups. |
| 6 | Unusual but benign incidents, e.g., trying to process a locked queue file. |
| 9 | Log internal queue id to external message id mappings. This can be useful for tracing a message as it travels between several hosts. |
| 12 | Several messages that are basically only of interest when debugging. |
| 16 | Verbose information regarding the queue. |

### C.7.4.4 File Modes

There are a number of files that can have a number of modes. The modes depend on what functionality you want and the level of security you require.

#### To Suid or Not to Suid?

Sendmail can safely be made setuid to root. At the point where it is about to exec a mailer, it checks to see if the userid is zero; if so, it resets the userid and groupid to a default (set by the u and g options). (This can be overridden by setting the S flag to the mailer for mailers that are trusted and must be called as root.) However, this causes mail processing to be accounted (using the sa command) to root rather than to the user sending the mail.

#### Temporary File Modes

The mode of all temporary files that sendmail creates is determined by the F option. Reasonable values for this option are 0600 and 0644. If the more permissive mode is selected, it is not necessary to run sendmail as root at all (even when running the queue).

#### Should My Alias Database be Writable?

Berkeley has the alias database (/usr/lib/aliases*) mode 666. There are some dangers inherent in this approach: any user can add himself to any list, or can "steal" any other user's mail. However, most users are trustworthy, and the cost of having a read-only database is greater than the expense of finding and eradicating the rare nasty person.

The database that sendmail actually used is represented by the two files aliases.dir and aliases.pag (both in /usr/lib). The mode on these files should match the mode on /usr/lib/aliases. If aliases is writable and the DBM files (aliases.dir and aliases.pag) are not, users will be unable to reflect their desired changes through to the actual database. However, if aliases is read-only and the DBM files are writable, a slightly sophisticated user can arrange to steal mail anyway.

If your DBM files are not writable by the world or you do not have auto-rebuild enabled (with the D option), you must be careful to reconstruct the alias database each time you change the text version:

```
newaliases
```

If this step is ignored or forgotten, any intended changes are also ignored or forgotten.

## C.7.5 Configuration File Details

This section describes the configuration file in detail, including hints on how to write one of your own if you have to.

There is one point that should be made clear immediately: the syntax of the configuration file is designed to be reasonably easy to parse, since this is done every time sendmail starts up, rather than easy for a human to read or write. On the future project list is a configuration-file compiler.

An overview of the configuration file is given first, followed by details of the semantics.

### C.7.5.1 The Syntax

The configuration file is organized as a series of lines, each of which begins with a single character defining the semantics for the rest of the line. Lines beginning with a space or a tab are continuation lines (although the semantics are not well defined in many places). Blank lines and lines beginning with a sharp symbol (#) are comments.

#### R and S - Rewriting Rules

The core of address parsing are the rewriting rules. These are an ordered production system. Sendmail scans through the set of rewriting rules looking for a match on the left hand side (LHS) of the rule. When a rule matches, the address is replaced by the right hand side (RHS) of the rule.

There are several sets of rewriting rules. Some of the rewriting sets are used internally and must have specific semantics. Other rewriting sets do not have specifically assigned semantics, and can be referenced by the mailer definitions or by other rewriting sets.

The syntax of these two commands are:

**S***n*

Sets the current ruleset being collected to n. If you begin a ruleset more than once it deletes the old definition.

**R***lhs rhs comments*

The fields must be separated by at least one tab character; there can be embedded spaces in the fields. The LHS is a pattern that is applied to the input. If it matches, the input is rewritten to the RHS. The comments are ignored.

### D – Define Macro

Macros are named with a single character. These can be selected from the entire ASCII set, but user-defined macros should be selected from the set of upper case letters only. Lower case letters and special symbols are used internally.

The syntax for macro definitions is:

D*xval*

where x is the name of the macro and val is the value it should have. Macros can be interpolated in most places using the escape sequence $x.

### C and F – Define Classes

Classes of words can be defined to match on the left hand side of rewriting rules. For example a class of all local names for this site might be created so that attempts to send to oneself can be eliminated. These can either be defined directly in the configuration file or read in from another file. Classes can be given names from the set of upper case letters. Lower case letters and special characters are reserved for system use.

The syntax is:

C*cword1 word2...*
F*cfile [ format ]*

The first form defines the class *c* to match any of the named words. It is permissible to split them among multiple lines; for example, the two forms:

```
CHmonet ucbmonet
```

and

```
CHmonet
CHucbmonet
```

are equivalent. The second form reads the elements of the class *c* from the named *file*; the format is a scanf pattern that should produce a single string.

### M - Define Mailer

Programs and interfaces to mailers are defined in this line. The format is:

M*name, {field=value}**

where *name* is the name of the mailer (used internally only) and the "field=name" pairs define attributes of the mailer. Fields are:

| | |
|---|---|
| Path | The pathname of the mailer |
| Flags | Special flags for this mailer |
| Sender | A rewriting set for sender addresses |
| Recipient | A rewriting set for recipient addresses |
| Argv | An argument vector to pass to this mailer |
| Eol | The end-of-line string for this mailer |
| Maxsize | The maximum message length to this mailer |

Only the first character of the field name is checked.

### H - Define Header

The format of the header lines that sendmail inserts into the message are defined by the *H* line. The syntax of this line is:

H*[?mflags?]hname: htemplate*

Continuation lines in this spec are reflected directly into the outgoing message. The htemplate is macro expanded before insertion into the message. If the *mflags* (surrounded by question marks) are specified, at least one of the specified flags must be stated in the mailer definition for this header to be automatically output. If one of these headers is in the input it is reflected to the output regardless of these flags.

Some headers have special semantics that are described below.

### O - Set Option

There are a number of random options that can be set from a configuration file. Options are represented by single characters. The syntax of this line is:

O*ovalue*

This sets option o to be value. Depending on the option, value can be a string, an integer, a boolean (with valid values *t, T, f, or F*; the default is TRUE), or a time interval.

### T – Define Trusted Users

Trusted users are those users who are permitted to override the sender address using the –f flag. These typically are *root*, *uucp*, and *network*, but on some users it may be convenient to extend this list to include other users, perhaps to support a separate UUCP login for each host. The syntax of this line is:

T*user1 user2...*

More than one of these lines can exist.

### P – Precedence Definitions

Values for the "Precedence:" field can be defined using the *P* control line. The syntax of this field is:

P*name=num*

When the name is found in a "Precedence:" field, the message class is set to num. Higher numbers mean higher precedence. Numbers less than zero have the special property that error messages are not returned. The default precedence is zero. For example, our list of precedences is:

```
Pfirst-class=0
Pspecial-delivery=100
Pjunk=-100
```

## C.7.5.2 The Semantics

This section describes the semantics of the configuration file.

### Special Macros, Conditionals

Macros are interpolated using the construct $*x*, where *x* is the name of the macro to be interpolated. In particular, lower case letters are reserved to have special semantics, used to pass information in or out of sendmail, and some special characters are reserved to provide conditionals, etc.

The following macros *must* be defined to transmit information into sendmail:

| | |
|---|---|
| e | The SMTP entry message |
| j | The official domain name for this site |
| l | The format of the UNIX from line |
| n | The name of the daemon (for error messages) |
| o | The set of operators in addresses |
| q | default format of sender address |

The $e macro is printed out when SMTP starts up. The first word must be the $j macro. The $j macro should be in RFC821 format. The $l and $n macros can be considered constants except under terribly unusual circumstances. The $o macro consists of a list of

characters that are considered tokens and that separate tokens when doing parsing. For example, if "r" were in the $o macro, the input address would be scanned as three tokens: "add," "r," and "ess." Finally, the $q macro specifies how an address should appear in a message when it is defaulted. For example, on our system these definitions are:

```
De$j Sendmail $v ready at $b
DnMAILER-DAEMON
DlFrom $g  $d
Do.:%@!^=/
Dq$g$?x ($x)$.
Dj$H.$D
```

An acceptable alternative for the $q macro is "$?x$x $.<$g>". These correspond to the following two formats:

```
eric@Berkeley (Eric Allman)
Eric Allman <eric@Berkeley>
```

Some macros are defined by sendmail for interpolation into argv's for mailers or for other contexts. These macros are:

| | |
|---|---|
| a | The origination date in Arpanet format |
| b | The current date in Arpanet format |
| c | The hop count |
| d | The date in UNIX (ctime) format |
| f | The sender (from) address |
| g | The sender address relative to the recipient |
| h | The recipient host |
| i | The queue id |
| p | Sendmail's pid |
| r | Protocol used |
| s | Sender's host name |
| t | A numeric representation of the current time |
| u | The recipient user |
| v | The version number of sendmail |
| w | The hostname of this site |
| x | The full name of the sender |
| y | The id of the sender's tty |
| z | The home directory of the recipient |

Three types of dates can be used. The $a and $b macros are in Arpanet format; $a is the time as extracted from the "Date:" line of the message (if there was one), and $b is the current date and time (used for postmarks). If no "Date:" line is found in the incoming message, $a is set to the current time also. The $d macro is equivalent to the $a macro in UNIX (ctime) format.

The $f macro is the id of the sender as originally determined; when mailing to a specific host the $g macro is set to the address of the sender *relative to the recipient*. For example, if I send to "bollard@matisse" from the machine "ucbarpa" the $f macro is "eric" and the $g macro is "eric@ucbarpa."

The $x macro is set to the full name of the sender. This can be determined in several ways. It can be passed as flag to sendmail. The second choice is the value of the "Full-name:" line in the header if it exists, and the third choice is the comment field of a "From:" line. If all of these fail, and if the message is being originated locally, the full name is looked up in the /etc/passwd file.

When sending, the $h, $u, and $z macros get set to the host, user, and home directory (if local) of the recipient. The first two are set from the $@ and $: part of the rewriting rules, respectively.

The $p and $t macros are used to create unique strings (e.g., for the "Message-Id:" field). The $i macro is set to the queue id on this host; if put into the timestamp line it can be extremely useful for tracking messages. The $y macro is set to the id of the terminal of the sender (if known); some systems like to put this in the Unix "From" line. The $v macro is set to be the version number of sendmail; this is normally put in timestamps and has been proven extremely useful for debugging. The $w macro is set to the name of this host if it can be determined. The $c field is set to the "hop count," i.e., the number of times this message has been processed. This can be determined by the -h flag on the command line or by counting the timestamps in the message.

The $r and $s fields are set to the protocol used to communicate with sendmail and the sending hostname; these are not supported in the current version.

Conditionals can be specified using the syntax:

```
$?x text1 $| text2 $.
```

This interpolates text1 if the macro $x is set, and text2 otherwise. The "else" ($|) clause can be omitted.

### Special Classes

The class $=w is set to be the set of all names this host is known by. This can be used to delete local hostnames.

### The Left Hand Side

The left hand side of rewriting rules contains a pattern. Normal words are simply matched directly. Metasyntax is introduced using a dollar sign. The metasymbols are:

| | |
|---|---|
| $* | Match zero or more tokens |
| $+ | Match one or more tokens |
| $– | Match exactly one token |
| $=*x* | Match any token in class *x* |
| $-*x* | Match any token not in class *x* |

If any of these match, they are assigned to the symbol $n for replacement on the right hand side, where n is the index in the LHS. For example, if the LHS:

```
$–:$+
```

is applied to the input:

```
UCBARPA:eric
```

the rule matches, and the values passed to the RHS are:

```
$1 UCBARPA
$2 eric
```

### The Right Hand Side

When the right hand side of a rewriting rule matches, the input is deleted and replaced by the right hand side. Tokens are copied directly from the RHS unless they are begin with a dollar sign. Metasymbols are:

| | |
|---|---|
| $*n* | Substitute indefinite token *n* from LHS |
| $>*n* " | Call" ruleset *n* |
| $#*mailer* | Resolve to *mailer* |
| $@*host* | Specify *host* |
| $:*user* | Specify *user* |

The $*n* syntax substitutes the corresponding value from a $+, $–, $*, $=, or $- match on the LHS. It can be used anywhere.

The $>*n* syntax causes the remainder of the line to be substituted as usual and then passed as the argument to ruleset *n*. The final value of ruleset *n* then becomes the substitution for this rule.

The $# syntax should *only* be used in ruleset zero. It causes evaluation of the ruleset to terminate immediately, and signals to sendmail that the address has completely resolved. The complete syntax is:

$#*mailer*$@*host*$:*user*

This specifies the {mailer, host, user} three-tuple necessary to direct the mailer. If the mailer is local the host part can be omitted. The mailer and host must be a single word, but the user can be multi-part.

A RHS can also be preceeded by a $@ or a $: to control evaluation. A $@ prefix causes the ruleset to return with the remainder of the RHS as the value. A $: prefix causes the rule to terminate immediately, but the ruleset to continue; this can be used to avoid continued application of a rule. The prefix is stripped before continuing.

The $@ and $: prefixes can preceed a $> spec; for example:

```
R$+     $:$>7$1
```

matches anything, passes that to ruleset seven, and continues; the $: is necessary to avoid an infinite loop.

### Semantics of Rewriting Rule Sets

There are five rewriting sets that have specific semantics. These are related as depicted below.

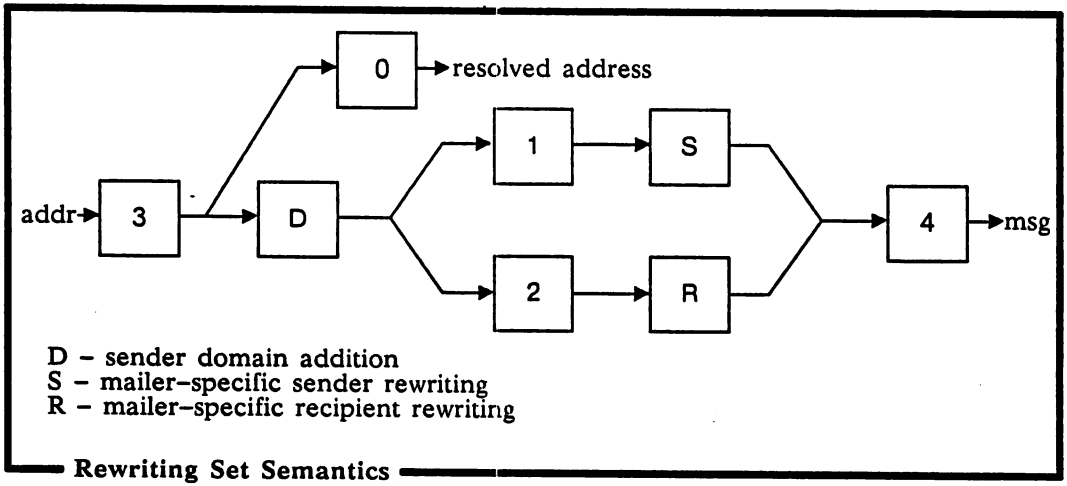


**Rewriting Set Semantics**

Ruleset three should turn the address into "canonical form." This form should have the basic syntax:

```
local-part@host-domain-spec
```

If no "@" sign is specified, the host-domain-spec can be appended from the sender address (if the C flag is set in the mailer definition corresponding to the sending mailer). Ruleset three is applied by sendmail before doing anything with any address.

Ruleset zero is applied after ruleset three to addresses that are going to actually specify recipients. It must resolve to a {*mailer*, *host*, *user*} triple. The *mailer* must be defined in the mailer definitions from the configuration file. The *host* is defined into the $*h* macro for use in the argv expansion of the specified mailer.

Rulesets one and two are applied to all sender and recipient addresses respectively. They are applied before any specification in the mailer definition. They must never resolve.

Ruleset four is applied to all addresses in the message. It is typically used to translate internal to external form.

### Mailer Flags Etc.

There are a number of flags that can be associated with each mailer, each identified by a letter of the alphabet. Many of them are assigned semantics internally. These are detailed in Section C.7.8. Any other flags can be used freely to conditionally assign headers to messages destined for particular mailers. –" . nr $1 error" –" . nr $2 mailer"

### The "Error" Mailer

The mailer with the special name "error" can be used to generate a user error. The (optional) host field is a numeric exit status to be returned, and the user field is a message to be printed. For example, the entry:

```
$#error$:Host unknown in this domain
```

on the RHS of a rule causes the specified error to be generated if the LHS matches. This mailer is only functional in ruleset zero.

## C.7.5.3 Building a Configuration File From Scratch

Building a configuration table from scratch is an extremely difficult job. Fortunately, it is almost never necessary to do so; nearly every situation that may come up can be resolved by changing an existing table. In any case, it is critical that you understand what it is that you are trying to do and come up with a philosophy for the configuration table. This section is intended to explain what the real purpose of a configuration table is and to give you some ideas for what your philosophy might be.

### What You are Trying to Do

The configuration table has three major purposes. The first and simplest is to set up the environment for sendmail. This involves setting the options, defining a few critical macros, etc., that are described in other places.

The second purpose is to rewrite addresses in the message. This should typically be done in two phases. The first phase maps addresses in any format into a canonical form. This should be done in ruleset three. The second phase maps this canonical form into the syntax appropriate for the receiving mailer. Sendmail does this in three subphases. Rulesets one and two are applied to all sender and recipient addresses respectively. After this, you can specify per-mailer rulesets for both sender and recipient addresses; this allows mailer-specific customization. Finally, ruleset four is applied to do any default conversion to external form.

The third purpose is to map addresses into the actual set of instructions necessary to get the message delivered. Ruleset zero must resolve to the internal form, which is in turn used as a pointer to a mailer descriptor. The mailer descriptor describes the interface requirements of the mailer.

## Philosophy

The particular philosophy you choose depends heavily on the size and structure of your organization. A few possible philosophies are presented here.

One general point applies to all of these philosophies: it is almost always a mistake to try to do full name resolution. For example, if you are trying to get names of the form "user@host" to the Arpanet, it does not pay to route them to

```
xyzvax!decvax!ucbvax!c70:user@host
```

since you then depend on several links not under your control. The best approach to this problem is to simply forward to "xyzvax!user@host" and let xyzvax worry about it from there. In summary, just get the message closer to the destination, rather than determining the full path.

## Large Site, Many Hosts – Minimum Information

Berkeley is an example of a large site, i.e., more than two or three hosts. We have decided that the only reasonable philosophy in our environment is to designate one host as the guru for our site. It must be able to resolve any piece of mail it receives. The other sites should have the minimum amount of information they can get away with. In addition, any information they do have should be hints rather than solid information.

For example, a typical site on our local ether network is "monet." Monet has a list of known ethernet hosts; if it receives mail for any of them, it can do direct delivery. If it receives mail for any unknown host, it just passes it directly to "ucbvax," our master host. Ucbvax may determine that the host name is invalid and reject the message, or may be able to do delivery. However, it is important to note that when a new ethernet host is added, the only host that *must* have its tables updated is ucbvax; the others *can* be updated as convenient, but this is not critical.

This picture is slightly muddied due to network connections that are not actually located on ucbvax. For example, our TCP connection is currently on "ucbarpa." However, monet *does not* know about this; the information is hidden totally between ucbvax and ucbarpa. Mail going from monet to a TCP host is transfered via the ethernet from monet to ucbvax, then via the ethernet from ucbvax to ucbarpa, and then is submitted to the Arpanet. Although this involves some extra hops, we feel this is an acceptable tradeoff.

An interesting point is that it would be possible to update monet to send TCP mail directly to ucbarpa if the load got too high; if monet failed to note a host as a TCP host it would go via ucbvax as before, and if monet incorrectly sent a message to ucbarpa it would still be sent by ucbarpa to ucbvax as before. The only problem that can occur is loops, as if ucbarpa thought that ucbvax had the TCP connection and vice versa. For this reason, updates should *always* happen to the master host first.

This philosophy results as much from the need to have a single source for the configuration files (typically built using m4 or some similar tool) as any logical need. Maintaining more than three separate tables by hand is essentially an impossible job.

### Small Site - Complete Information

A small site (two or three hosts) may find it more reasonable to have complete information at each host. This would require that each host know exactly where each network connection is, possibly including the names of each host on that network. As long as the site remains small and the the configuration remains relatively static, the update problem will probably not be too great.

### Single Host

This is in some sense the trivial case. The only major issue is trying to insure that you don't have to know too much about your environment. For example, if you have a UUCP connection you might find it useful to know about the names of hosts connected directly to you, but this is really not necessary since this can be determined from the syntax.

### Relevant Issues

The canonical form you use should almost certainly be as specified in the Arpanet protocols RFC819 and RFC822. Copies of these RFC's are included on the sendmail tape as doc/rfc819.lpr and doc/rfc822.lpr.

RFC822 describes the format of the mail message itself. Sendmail follows this RFC closely, to the extent that many of the standards described in this document can not be changed without changing the code. In particular, the following characters have special interpretations:

```
< > ( ) " \
```

Any attempt to use these characters for other than their RFC822 purpose in addresses is probably doomed to disaster.

RFC819 describes the specifics of the domain-based addressing. This is touched on in RFC822 as well. Essentially each host is given a name that is a right-to-left dot qualified pseudo-path from a distinguished root. The elements of the path need not be physical hosts; the domain is logical rather than physical. For example, at Berkeley one valid host is "a.cc.berkeley.arpa"; reading from right to left, "arpa" is a top level domain (related to, but not limited to, the physical Arpanet), "berkeley" is both an Arpanet host and a logical domain that is actually interpreted by a host called ucbvax (which is actually just the "major" host for this domain), "cc" represents the Computer Center, (in this case a strictly logical entity), and "a" is a host in the Computer Center; this particular host happens to be connected via berknet, but other hosts might be connected via one of two ethernets or some other network.

Beware when reading RFC819 that there are a number of errors in it.

### How to Proceed

Once you have decided on a philosophy, it is worth examining the available configuration tables to decide if any of them are close enough to borrow major parts of. Even under the worst of conditions, there is a fair amount of boiler plate that can be collected safely.

The next step is to build ruleset three. This is the hardest part of the job. Beware of doing too much to the address in this ruleset, since anything you do reflects through to the message. In particular, stripping of local domains is best deferred, since this can leave you with addresses with no domain spec at all. Since sendmail likes to append the sending domain to addresses with no domain, this can change the semantics of addresses. Also try to avoid fully qualifying domains in this ruleset. Although technically valid, this can lead to unpleasantly and unnecessarily long addresses reflected into messages. The Berkeley configuration files define ruleset nine to qualify domain names and strip local domains. This is called from ruleset zero to get all addresses into a cleaner form.

Once you have ruleset three finished, the other rulesets should be relatively trivial. If you need hints, examine the supplied configuration tables.

### Testing the Rewriting Rules - the -bt Flag

When you build a configuration table, you can do a certain amount of testing using the test mode of sendmail. For example, you could invoke sendmail as:

```
sendmail -bt -Ctest.cf
```

which would read the configuration file test.cf and enter test mode. In this mode, you enter lines of the form:

```
rwset address
```

where *rwset* is the rewriting set you want to use and *address* is an address to apply the set to. Test mode shows you the steps it takes as it proceeds, finally showing you the address it ends up with. You can use a comma separated list of rwsets for sequential application of rules to an input; ruleset three is always applied first. For example:

```
1,21,4 monet:bollard
```

first applies ruleset three to the input "monet:bollard." Ruleset one is then applied to the output of ruleset three, followed similarly by rulesets twenty-one and four.

If you need more detail, you can also use the -d21 flag to turn on more debugging. For example,

```
sendmail -bt -d21.99
```

turns on an incredible amount of information; a single word address is probably going to print out several pages worth of information.

### Building Mailer Descriptions

To add an outgoing mailer to your mail system, you must define the characteristics of the mailer.

Each mailer must have an internal name. This can be arbitrary, except that the names "local" and "prog" must be defined.

The pathname of the mailer must be given in the P field. If this mailer should be accessed via an IPC connection, use the string "[IPC]" instead.

The F field defines the mailer flags. Specify an -f or -r flag to pass the name of the sender. These flags are only passed if they were passed to sendmail, so that mailers that give errors under some circumstances can be placated. If the mailer is not picky you can just specify -f $g in the argv template. If the mailer must be called as *root* the -S flag should be given; this does not reset the userid before calling the mailer*. If this mailer is local (i.e., performs final delivery rather than another network hop) the l flag should be given. Quote characters (backslashes and double quotes) can be stripped from addresses if the s flag is specified; if this is not given they are passed through. If the mailer is capable of sending to more than one user on the same host in a single transaction the m flag should be stated. If this flag is on, the argv template containing $*u* is repeated for each unique user on a given host. The e flag marks the mailer as being "expensive," which causes sendmail to defer connection until a queue run**.

An unusual case is the C flag. This flag applies to the mailer that the message is received from, rather than the mailer being sent to; if set, the domain spec of the sender (i.e., the "@host.domain" part) is saved and is appended to any addresses in the message that do not already contain a domain spec. For example, a message of the form:

```
From: eric@ucbarpa
To: wnj@monet, mckusick
```

is modified to:

```
From: eric@ucbarpa
To: wnj@monet, mckusick@ucbarpa
```

*if and only if* the C flag is defined in the mailer corresponding to "eric@ucbarpa."

Other flags are described in Section C.7.8.

---

* Sendmail must be running setuid to root for this to work.

** The "c" configuration option must be given for this to be effective.

The S and R fields in the mailer description are permailer rewriting sets to be applied to sender and recipient addresses respectively. These are applied after the sending domain is appended and the general rewriting sets (numbers one and two) are applied, but before the output rewrite (ruleset four) is applied. A typical use is to append the current domain to addresses that do not already have a domain. For example, a header of the form:

```
From: eric
```

might be changed to be:

```
From: eric@ucbarpa
```

or

```
From: ucbvax!eric
```

depending on the domain it is being shipped into. These sets can also be used to do special purpose output rewriting in cooperation with ruleset four.

The E field defines the string to use as an end-of-line indication. A string containing only newline is the default. The usual backslash escapes (\r, \n, \f, \b) can be used.

Finally, an argv template is given as the E field. It can have embedded spaces. If there is no argv with a *$u* macro in it, sendmail speaks SMTP to the mailer. If the pathname for this mailer is "[IPC]," the argv should be

```
IPC $h [ port ]
```

where port is the optional port number to connect to.

For example, the specifications:

```
Mlocal, P=/bin/mail, F=rlsm S=10, R=20, A=mail -d $u
Mether, P=[IPC], F=meC, S=11, R=21, A=IPC $h, M=100000
```

specifies a mailer to do local delivery and a mailer for ethernet delivery. The first is called "local," is located in the file /bin/mail, takes a picky –r flag, does local delivery, quotes should be stripped from addresses, and multiple users can be delivered at once; ruleset ten should be applied to sender addresses in the message and ruleset twenty should be applied to recipient addresses; the argv to send to a message is the word "mail," the word "–d," and words containing the name of the receiving user. If an –r flag is inserted, it is between the words "mail" and "–d." The second mailer is called "ether," it should be connected to via an IPC connection, it can handle multiple users at once, connections should be deferred, and any domain from the sender address should be appended to any receiver name without a domain; sender addresses should be processed by ruleset eleven and recipient addresses by ruleset twenty-one. There is a 100,000 byte limit on messages passed through this mailer.

## C.7.6 Command Line Arguments

Arguments must be presented with flags before addresses. The flags are:

-f *addr* The sender's machine address is *addr*. This flag is ignored unless the real user is listed as a "trusted user" or if *addr* contains an exclamation point (because of certain restrictions in UUCP).

-r *addr* An obsolete form of *-f*.

-h *cnt* Sets the "hop count" to *cnt*. This represents the number of times this message has been processed by sendmail (to the extent that it is supported by the underlying networks). *Cnt* is incremented during processing, and if it reaches MAXHOP (currently 30) sendmail throws away the message with an error.

-F*name* Sets the full name of this user to *name*.

-n Don't do aliasing or forwarding.

-t Read the header for "To:", "Cc:", and "Bcc:" lines, and send to everyone listed in those lists. The "Bcc:" line is deleted before sending. Any addresses in the argument vector is deleted from the send list.

-b*x* Set operation mode to *x*. Operation modes are:

| | |
|---|---|
| m | Deliver mail (default) |
| a | Run in arpanet mode (see below) |
| s | Speak SMTP on input side |
| d | Run as a daemon |
| t | Run in test mode |
| v | Just verify addresses, don't collect or deliver |
| i | Initialize the alias database |
| p | Print the mail queue |
| z | Freeze the configuration file |

The special processing for the ARPANET includes reading the "From:" line from the header to find the sender, printing ARPANET style messages (preceded by three digit reply codes for compatibility with the FTP protocol [Neigus73, Postel74, Postel77]), and ending lines of error messages with <CRLF>.

-q*time* Try to process the queued up mail. If the time is given, a sendmail runs through the queue at the specified interval to deliver queued mail; otherwise, it only runs once.

-C*file* Use a different configuration file.

-d*level* Set debugging level.

-o*xvalue* Set option *x* to the specified *value*. These options are described in Section C.7.7.

There are a number of options that can be specified as primitive flags (provided for compatibility with delivermail). These are the e, i, m, and v options. Also, the f option can be specified as the -s flag.

## C.7.7 Configuration Options

The following options can be set using the –o flag on the command line or the O line in the configuration file:

A*file* Use the named *file* as the alias file. If no file is specified, use *aliases* in the current directory.

a If set, wait for an "@:@" entry to exist in the alias database before starting up. If it does not appear in five minutes, rebuild the database.

c If an outgoing mailer is marked as being expensive, don't connect immediately. This requires that queueing be compiled in, since it depends on a queue run process to actually send the mail.

d*x* Deliver in mode *x*. Valid modes are:

| | |
|---|---|
| i | Deliver interactively (synchronously) |
| b | Deliver in background (asynchronously) |
| q | Just queue the message (deliver during queue run) |

D If set, rebuild the alias database if necessary and possible. If this option is not set, sendmail never rebuilds the alias database unless explicitly requested using *–bi*.

e*x* Dispose of errors using mode *x*. The values for *x* are:

| | |
|---|---|
| p | Print error messages (default) |
| q | No messages, just give exit status |
| m | Mail back errors |
| w | Write back errors (mail if user not logged in) |
| e | Mail back errors and give zero exit stat always |

F*n* The temporary file mode, in octal. 644 and 600 are good choices.

f Save Unix-style "From" lines at the front of headers. Normally they are assumed redundant and discarded.

g*n* Set the default group id for mailers to run in to *n*.

H*file* Specify the help file for SMTP.

i Ignore dots in incoming messages.

L*n* Set the default log level to *n*.

M*xvalue* Set the macro *x* to *value*. This is intended only for use from the command line.

m Send to me too, even if I am in an alias expansion.

o Assume that the headers may be in old format, i.e., spaces delimit names. This actually turns on an adaptive algorithm: if any recipient address con-

tains a comma, parenthesis, or angle bracket, it is assumed that commas already exist. If this flag is not on, only commas delimit names. Headers are always output with commas between the names.

Q*dir* Use the named *dir* as the queue directory.

r*time* Timeout reads after *time* interval.

S*file* Log statistics in the named *file*.

s Be super-safe when running things, i.e., always instantiate the queue file, even if you are going to attempt immediate delivery. Sendmail always instantiates the queue file before returning control the the client under any circumstances.

T*time* Set the queue timeout to *time*. After this interval, messages that have not been successfully sent are returned to the sender.

t*S,D* Set the local timezone name to *S* for standard time and *D* for daylight time.

u*n* Set the default userid for mailers to *n*. Mailers without the *S* flag in the mailer definition run as this user.

v Run in verbose mode.

## C.7.8 Mailer Flags

The following flags can be set in the mailer description.

f The mailer wants a *–f from* flag, but only if this is a network forward operation (i.e., the mailer gives an error if the executing user does not have special permissions).

r Same as *f*, but sends a *–r* flag.

S Don't reset the userid before calling the mailer. This would be used in a secure environment where *sendmail* ran as root. This could be used to avoid forged addresses. This flag is suppressed if given from an "unsafe" environment (e.g, a user's mail.cf file).

n Do not insert a UNIX-style "From" line on the front of the message.

l This mailer is local (i.e., final delivery is performed).

s Strip quote characters off of the address before calling the mailer.

m This mailer can send to multiple users on the same host in one transaction. When a *$u* macro occurs in the *argv* part of the mailer definition, that field is repeated as necessary for all qualifying users.

F This mailer wants a "From:" header line.

D This mailer wants a "Date:" header line.

M This mailer wants a "Message-Id:" header line.

x This mailer wants a "Full-Name:" header line.

P This mailer wants a "Return-Path:" line.

u Preserve upper case in user names for this mailer.

h Preserve upper case in host names for this mailer.

A This is an Arpanet-compatible mailer; set appropriate modes.

U This mailer wants Unix-style "From" lines with the ugly UUCP-style "remote from <host>" on the end.

e This mailer is expensive to connect to, so try to avoid connecting normally; any necessary connection occurs during a queue run.

X This mailer want to use the hidden dot algorithm as specified in RFC821; basically, any line beginning with a dot has an extra dot prepended (to be stripped at the other end). This insures that lines in the message containing a dot do not terminate the message prematurely.

L Limit the line lengths as specified in RFC821.

P Use the return-path in the SMTP "MAIL FROM:" command rather than just the return address; although this is required in RFC821, many hosts do not process return paths properly.

I This mailer will be speaking SMTP to another sendmail – as such it can use special protocol features. This option is not required (i.e., if this option is omitted the transmission still operates successfully, although perhaps not as efficiently as possible).

C If mail is *received* from a mailer with this flag set, any addresses in the header that do not have an at sign ("@") after being rewritten by ruleset three, have the "@domain" clause from the sender tacked on. This allows mail with headers of the form:

```
From: usera@hosta
To: userb@hostb, userc
```

to be rewritten as:

```
From: usera@hosta
To: userb@hostb, userc@hosta
```

automatically.

## C.7.9 Other Configuration Changes

Some configuration changes can be made by recompiling sendmail. These are located in three places:

md/config.m4 These contain operating-system dependent descriptions. They are interpolated into the Makefiles in the src and aux directories. This includes information about what version of UNIX you are running, what libraries you have to include, etc.

src/conf.h Configuration parameters that can be tweaked by the installer are included in conf.h.

src/conf.c Some special routines and a few variables can be defined in conf.c. For the most part these are selected from the settings in conf.h.

### C.7.9.1 Parameters in md/config.m4

The following compilation flags can be defined in the m4CONFIG macro in md/config.m4 to define the environment in which you are operating.

V6 If set, this compiles a system with eight-bit user id's, single character tty id's, etc.

VMUNIX If set, you are assumed to have a Berkeley 4BSD or 4.1BSD, including the vfork system call, special types defined in <sys/types.h> (e.g, u_char), etc.

If none of these flags are set, a version 7 system is assumed.

You also must specify what libraries to link with sendmail in the m4LIBS macro. Most notably, you must include if you are running a 4.1BSD system.

### C.7.9.2 Parameters in src/conf.h

Parameters and compilation options are defined in conf.h. Most of these need not normally be tweaked; common parameters are all in sendmail.cf. However, the sizes of certain primitive vectors, etc., are included in this file. The numbers following the parameters are their default value.

MAXLINE [256]
The maximum line length of any input line. If message lines exceed this length they are still processed correctly; however, header lines, configuration file lines, alias lines, etc., must fit within this limit.

MAXNAME [128]
The maximum length of any name, such as a host or a user name.

MAXFIELD [2500]
The maximum total length of any header field, including continuation lines.

MAXPV [40]
The maximum number of parameters to any mailer. This limits the number of recipients that can be passed in one transaction.

MAXHOP [30]
When a message has been processed more than this number of times, sendmail rejects the message on the assumption that there has been an aliasing loop. This can be determined from the –h flag or by counting the number of trace fields (i.e, "Received:" lines) in the message header.

MAXATOM [100]
The maximum number of atoms (tokens) in a single address. For example, the address "eric@Berkeley" is three atoms.

MAXMAILERS [25]
The maximum number of mailers that can be defined in the configuration file.

MAXRWSETS [30]
The maximum number of rewriting sets that can be defined.

MAXPRIORITIES [25]
The maximum number of values for the "Precedence:" field that can be defined (using the *P* line in sendmail.cf).

MAXTRUST [30]
The maximum number of trusted users that can be defined (using the *T* line in sendmail.cf).

A number of other compilation options exist. These specify whether or not specific code should be compiled in.

DBM — If set, the DBM package in UNIX is used (see the *Concentrix-2800 System Reference:System Calls and C Library Subroutines:* dbm). If not set, a much less efficient algorithm for processing aliases is used.

DEBUG — If set, debugging information is compiled in. To actually get the debugging output, use the –d flag.

LOG — If set, the syslog routine in use at some sites is used. This makes an informational log record for each message processed, and makes a higher priority log record for internal system errors.

QUEUE — This flag should be set to compile in the queueing code. If this is not set, mailers must accept the mail immediately or it is returned to the sender.

SMTP — If set, the code to handle user and server SMTP is compiled in. This is only necessary if your machine has some mailer that speaks SMTP.

DAEMON — If set, code to run a daemon is compiled in. This code is for 4.2BSD if the NVMUNIX flag is specified; otherwise, 4.1a BSD code is used. Beware however that there are bugs in the 4.1a code that make it impossible for sendmail to work correctly under heavy load.

UGLYUUCP If you have a UUCP host adjacent to you that is not running a reasonable version of rmail, you must set this flag to include the "remote from sysname" info on the from line. Otherwise, UUCP gets confused about where the mail came from.

NOTUNIX If you are using a non-UNIX mail format, you can set this flag to turn off special processing of UNIX-style "From " lines.

### C.7.9.3 Configuration in src/conf.c

Not all header semantics are defined in the configuration file. Header lines that should only be included by certain mailers (as well as other more obscure semantics) must be specified in the *HdrInfo* table in conf.c. This table contains the header name (which should be in all lower case) and a set of header control flags (described below), The flags are:

H_ACHECK Normally when the check is made to see if a header line is compatible with a mailer, sendmail does not delete an existing line. If this flag is set, sendmail deletes even existing header lines. That is, if this bit is set and the mailer does not have flag bits set that intersect with the required mailer flags in the header definition in sendmail.cf, the header line is *always* deleted.

H_EOH If this header field is set, treat it like a blank line, i.e., it signals the end of the header and the beginning of the message text.

H_FORCE Add this header entry even if one existed in the message before. If a header entry does not have this bit set, sendmail does not add another header line if a header line of this name already existed. This would normally be used to stamp the message by everyone who handled it.

H_TRACE If set, this is a timestamp (trace) field. If the number of trace fields in a message exceeds a preset amount the message is returned on the assumption that it has an aliasing loop.

H_RCPT If set, this field contains recipient addresses. This is used by the -t flag to determine who to send to when it is collecting recipients from the message.

H_FROM This flag indicates that this field specifies a sender. The order of these fields in the *HdrInfo* table specifies sendmail's preference for which field to return error messages to.

Let's look at a sample *HdrInfo* specification:

```
struct hdrinfo HdrInfo[] =
{
        /* originator fields, most to least significant */
    "resent-sender",    H_FROM,
    "resent-from",      H_FROM,
    "sender",       H_FROM,
    "from",         H_FROM,
    "full-name",        H_ACHECK,
        /* destination fields */
    "to",           H_RCPT,
    "resent-to",        H_RCPT,
    "cc",           H_RCPT,
        /* message identification and control */
    "message",      H_EOH,
    "text",             H_EOH,
        /* trace fields */
    "received",     H_TRACE|H_FORCE,

    NULL,           0,
};
```

This structure indicates that the "To:", "Resent-To:", and "Cc:" fields all specify recipient addresses. Any "FullName:" field is deleted unless the required mailer flag (indicated in the configuration file) is specified. The "Message:" and "Text:" fields terminate the header; these are specified in new protocols [NBS80] or used by random dissenters around the network world. The "Received:" field is always added, and can be used to trace messages.

There are a number of important points here. First, header fields are not added automatically just because they are in the *HdrInfo* structure; they must be specified in the configuration file in order to be added to the message. Any header fields mentioned in the configuration file but not mentioned in the *HdrInfo* structure have default processing performed; that is, they are added unless they were in the message already. Second, the *HdrInfo* structure only specifies cliched processing; certain headers are processed specially by ad hoc code regardless of the status specified in *HdrInfo*. For example, the "Sender:" and "From:" fields are always scanned on ARPANET mail to determine the sender; this is used to perform the "return to sender" function. The "From:" and "FullName:" fields are used to determine the full name of the sender if possible; this is stored in the macro $x and used in a number of ways.

The file conf.c also contains the specification of ARPANET reply codes. There are four classifications these fall into:

```
char   Arpa_Info[] =       "050";       /* arbitrary info */
char   Arpa_TSyserr[] =    "455";       /* some (transient)
system error */
char   Arpa_PSyserr[] =    "554"; /* some (transient) system
error */
char   Arpa_Usrerr[] =     "554"; /* some (fatal) user error */
```

The class *Arpa_Info* is for any information that is not required by the protocol, such as forwarding information. *Arpa_TSyserr* and *Arpa_PSyserr* is printed by the syserr routine. TSyserr is printed out for transient errors, whereas PSyserr is printed for permanent errors; the distinction is made based on the value of errno. Finally, *Arpa_Usrerr* is the result of a user error and is generated by the usrerr routine; these are generated when the user has specified something wrong, and hence the error is permanent, i.e., it does not work simply by resubmitting the request.

If it is necessary to restrict mail through a relay, the checkcompat routine can be modified. This routine is called for every recipient address. It can return *TRUE* to indicate that the address is acceptable and mail processing continues, or it can return *FALSE* to reject the recipient. If it returns false, it is up to checkcompat to print an error message (using usrerr) saying why the message is rejected. For example, checkcompat could read:

```
bool
checkcompat(to)
        register ADDRESS *to;
{
        if (MsgSize > 50000 && to->q_mailer != LocalMailer)
        {
                usrerr("Message too large for non-local deliv-
ery");
                NoReturn = TRUE;
                return (FALSE);
        }
        return (TRUE);
}
```

This would reject messages greater than 50000 bytes unless they were local. The *NoReturn* flag can be sent to supress the return of the actual body of the message in the error return. The actual use of this routine is highly dependent on the implementation, and use should be limited.

## C.7.10 Summary of Support Files

This is a summary of the support files that sendmail creates or generates.

| | |
|---|---|
| /usr/lib/sendmail | The binary of sendmail. |
| /usr/bin/newaliases | A link to /usr/lib/sendmail; causes the alias database to be rebuilt. Running this program is completely equivalent to giving sendmail the -bi flag. |
| /usr/bin/mailq | Prints a listing of the mail queue. This program is equivalent to using the -bp flag to sendmail. |
| /usr/lib/sendmail.cf | The configuration file, in textual form. |
| /usr/lib/sendmail.hf | The SMTP help file. |
| /usr/lib/sendmail.st | A statistics file; need not be present. |

| | |
|---|---|
| /usr/lib/aliases | The textual version of the alias file. |
| /usr/lib/aliases.{pag,dir} | The alias file in dbm format. |
| /etc/syslog | The program to do logging. |
| /etc/syslog.conf | The configuration file for syslog. |
| /etc/syslog.pid | Contains the process id of the currently running syslog. |
| /usr/spool/mqueue | The directory in which the mail queue and temporary files reside. |
| /usr/spool/mqueue/qf* | Control (queue) files for messages. |
| /usr/spool/mqueue/df* | Data files. |
| /usr/spool/mqueue/lf* | Lock files. |
| /usr/spool/mqueue/tf* | Temporary versions of the qf files, used during queue file rebuild. |
| /usr/spool/mqueue/nf* | A file used when creating a unique id. |
| /usr/spool/mqueue/xf* | A transcript of the current session. |

# APPENDIX D
# DISK MAINTENANCE AND INSTALL UTILITIES

The disk maintenance utilities run under a Diagnostix-2800 or Concentrix-2800 operating system that accesses RAM disk (physical memory simulates a disk device), and permit you to execute certain utilities without a root disk (or with a bad root disk).

## D.1 CONTENTS OF RAM DISKS

The RAM disks contain the disk maintenance and install utilities for Concentrix-2800 and Diagnostix-2800.

### D.1.1 Concentrix-2800 RAM Disk

The Concentrix-2800 RAM disk contains the following utilities:

| | |
|---|---|
| /bin/adb | assembler-level debugger |
| /bin/afio | tape archiver |
| /bin/bfio | tape archiver |
| /bin/cat | display text file |
| /bin/cp | copy files, directories, subtrees |
| /bin/csh | C shell |
| /bin/dd | convert and copy a file |
| /bin/df | report free disk space on file systems |
| /bin/echo | echo arguments |
| /bin/findroot | find root and swap partitions during installation |
| /bin/grep | search a file for a pattern |
| /bin/iu.cx.sh | install or upgrade root partition |
| /bin/ln | make links |
| /bin/ls | list contents of directory |
| /bin/make | maintain program groups |
| /bin/mkdir | create directory |
| /bin/mt | magnetic tape manipulating program |
| /bin/pickroot | allow user to pick root partition during installation |
| /bin/rm | remove (unlink) files |
| /bin/sh | Bourne shell |
| /bin/stty | set terminal options |
| /bin/sync | update the super block |
| /bin/tar | tape archiver |
| /bin/test | evaluate expression and return status |
| /etc/diskf | disk utility program |
| /etc/disktab | disk description file |
| /etc/dlabel | generate disk labels from prototype file |

| | |
|---|---|
| /etc/dmesg | collect system diagnostic messages in error log |
| /etc/fastboot | reboot system without checking disks |
| /etc/fasthalt | halt system in console mode without checking disks |
| /etc/fsck | file system consistency check / interactive repair * |
| /etc/halt | halt system in console mode |
| /etc/init | system initialization and process control |
| /etc/makefs | construct a file system |
| /etc/mknod | build a special file |
| /etc/mount | mount file system |
| /etc/mtab | mounted file system table |
| /etc/reboot | reboot system |
| /etc/restore | incremental file system restore |
| /etc/shutdown | shut down system at a given time |
| /etc/umount | dismount file systems |
| /dev/mkdisknod | make special files for standard disk devices |
| /dev/mknewdevs | make default devices for initial installation |
| /dev/mkptynod | make special files for pseudo terminals |
| /dev/mkttynod | make special files for terminals |

**Warning:** Do not use fsck on unformatted, unlabelled root disks. Enter the following command to see if the root disk is labelled and formatted:

```
/etc/diskf /dev/rxdv00z S
```

## D.1.2 Diagnostix-2800 Maintenance Utilities RAM Disk

The Diagnostix-2800 Maintenance Utilities RAM disk contains Diagnostix-2800 executables as follows:

| | |
|---|---|
| /bin/bfio | tape archiver |
| /bin/cat | display text file |
| /bin/cp | copy files, directories, subtrees |
| /bin/dd | convert and copy a file |
| /bin/ls | list contents of directory |
| /bin/mkdir | create directory |
| /bin/ls | list contents of directory |
| /bin/mt | magnetic tape manipulating program |
| /bin/rm | remove (unlink) files |
| /bin/sync | update the superblock |
| /bin/tar | tape archiver |
| /etc/diskf | disk utility program |
| /etc/disktab | disk description file |
| /etc/dlabel | generate disk labels from prototype file |
| /etc/fsck | file system consistency check / interactive repair * |
| /etc/makefs | construct a file system |
| /etc/mknod | build a special file |
| /etc/mount | mount file system |
| /etc/reboot | reboot system |
| /etc/umount | mount file system |

**Warning:** Do not use fsck on unformatted, unlabelled root disks. Enter the following command to see if the root disk is labelled and formatted:

```
/etc/diskf /dev/rxdv00z S
```

### D.1.3 Diagnostix-2800 Install Utilities RAM Disk

The Diagnostix-2800 Install Utilities RAM disk contains the following utilities:

| | |
|---|---|
| /bin/bfio | tape archiver |
| /bin/cat | display text file |
| /bin/dd | convert and copy a file |
| /bin/ls | list contents of directory |
| /bin/mkdir | create directory |
| /bin/mt | magnetic tape manipulating program |
| /bin/rm | remove (unlink) files |
| /bin/sync | update the superblock |
| /etc/diskf | disk utility program |
| /etc/disktab | disk description file |
| /etc/dlabel | generate disk labels from prototype file |
| /etc/makefs | construct a file system |
| /etc/mknod | build a special file |
| /etc/mount | mount file system |
| /etc/reboot | reboot system |
| /etc/restore | incremental file system restore |
| /etc/umount | mount file system |

## D.2 PROCEDURES FOR USING RAM DISK UTILITIES

The procedures for using the Concentrix-2800 and Diagnostix-2800 RAM Disk Utilities follow:

### D.2.1 Concentrix-2800 RAM Disk Utilities

1. If the system is running in multi-user Concentrix-2800, halt the system in console mode with the shutdown command:

   ```
   # /etc/shutdown -H +5 fix root disk RETURN
   ```

   If the system is running in single-user Concentrix-2800, halt the system in console mode with the reboot command:

   ```
   # /etc/reboot -H RETURN
   ```

   If the system is powered down or the operating system is not running, power the system on (or off and on) using the power switch on the front panel. Turn the power switch to the ON (12:00) position, and press the RESET button. This stops the system in console mode.

2. Type o, S RETURN, u (as shown below) at the console prompts to start the system and stop in Diagnostix-2800 mode (you can type ? for help at any of the prompts in the steps below):

   ```
   cons>Options =  Enter RETURN to keep, SPACE to clear, or new op-
   tions.
   New Options? S RETURN

   cons>u
   ```

3. Enter CTRL d at the dunix-sh# prompt.

   ```
   dunix-sh# CTRL d
   ```

   The following messages appear:

   ```
   FX/2800 Startup

   Entering Console Boot Loop.
   ```

4. At the series of questions about reboot with flags, enter R to access the RAM disk, e to enable booting from the RAM disk, and y to boot Concentrix-2800, followed by RETURN each time, as shown below:

   ```
   Reboot with flags set to " "? (y for yes, ? for help): R RETURN

   Ramdisk booting is currently disabled.
   Do you want to enable (e) or disable (d)?e RETURN
   Reboot with flags set to " "? (y for yes, ? for help): y RETURN
   ```

5. When the Concentrix-2800 Utilities Menu appears, select the desired option and press RETURN. To use the Concentrix-2800 RAM disk utilities, choose option X:

**Concentrix Utilities Menu**

```
                    Concentrix Utilities Menu

        (Please Choose one of the following options and hit <return>.)

Option  Action                                          Default Command
------- ----------------------------------------------- -------------------
   I       Create root (/) and /usr file systems and
           Install Concentrix                           (/bin/iu.cx.sh i)
   U       Upgrade Concentrix operating system
           Upgrades both root (/) and /usr file systems (/bin/iu.cx.sh u)
   NR      Nonstandard upgrade of Concentrix root (/)
           partition only                               (/bin/iu.cx.sh nr)
   NU      Nonstandard upgrade of Concentrix /usr
           partition only                               (/bin/iu.cx.sh nu)
   R       Boot to Multi-User Concentrix                (/etc/reboot)
   s       Boot to Single-User Concentrix               (/etc/reboot -s)
   S       Boot to Diagnostix                           (/etc/reboot -S)
   H       Boot to cons> prompt                         (/etc/reboot -H)
   X       Exit to Concentrix prompt.

Option? > X RETURN
```

The following message appears, followed by the Concentrix Ramdisk prompt. The Concentrix-2800 RAM disk commands are available at this point.

```
Exiting Concentrix Install Utilities Menu

Concentrix Ramdisk#
```

To exit from the Concentrix-2800 RAM disk utilities, type CTRL D, then one of the following options when the menu appears: R, s, S, H.

## D.2.2 Diagnostix-2800 RAM Disk Maintenance Utilities

1. Insert the Concentrix-2800 distribution tape in the Exabyte drive.

2. If the system is running in multi-user Concentrix-2800, halt the system in console mode with the shutdown command:

```
# /etc/shutdown -H +5 fix root disk RETURN
```

   If the system is running in single-user Concentrix-2800, halt the system in console mode with the reboot command:

```
# /etc/reboot -H RETURN
```

   If the system is powered down or the operating system is not running, power the system on (or off and on) using the power switch on the front panel. Turn the power switch to the ON (12:00) position, and press the RESET button. This stops the system in console mode.

3. Type o, B RETURN (as shown below) at the console prompts to start the system and stop in VMEbus booter mode (you can type ? for help at any of the prompts in the steps below):

```
cons>Options = Enter RETURN to keep, SPACE to clear, or new op-
tions.
New Options? B RETURN
```

4. Select the tape drive as the device containing the booter program:

```
cons>device boot from
(d)isk (t)ape = t
```

   Note: You only need to type d and t. The system fills in the rest of the input.

5. Select the default controller for the tape drive by pressing RETURN:

```
controller (0-7) [return for default] = RETURN
```

6. The system loads the booter. After the message announcing the Alliant VMEbus booter appears, type %r RETURN to direct the booter to copy files onto the RAM disk, as shown below:

```
Alliant VMEbus Booter [made: Wed May 9 19:32:23 EDT 1990]
Type %h for help

boot from:  xdv(0,0,p)dunix: %r RETURN
```

7. Direct the booter to load the RAM disk from the first file on the tape, as shown below:

```
load ramdisk from:  xdv(0,0,p)ramdisk.init: hs(1) RETURN
```

8. Direct the booter to load Diagnostix-2800 from the third file on the tape, as shown below:

```
boot from:  hs(1)dunix: hs(3) RETURN
```

9. Direct the booter to load the Diagnostix-2800 dddt program from the next file on the tape, as shown below:

```
boot from:  hs(3)dddt:  hs(*) RETURN
```

10. When the Diagnostix-2800 Disk Utilities Menu appears, select the desired option and press RETURN. To use the Diagnostix-2800 RAM disk utilities, choose option X:

**Disk Utilities Menu**

```
                    Disk Utilities Menu

      (Please Choose one of the following options and hit <return>.)

Option  Action                                          Default Command
------- ----------------------------------------------- -----------------
   S    Boot to Diagnostix.                             (/etc/reboot -S)
   H    Boot to cons> prompt.                           (/etc/reboot -H)
   X    Exit to Diagnostix prompt.

Option? > X RETURN
```

The following message appears, followed by the Diagnostix Ramdisk prompt.

```
        Exiting Diagnostix Disk Utilities Menu

Diagnostix Ramdisk#
```

To exit from the Diagnostix-2800 RAM disk maintenance utilities, type CTRL D, then one of the following options when the menu appears: S, H.

## D.2.3 Diagnostix-2800 RAM Disk Install Utilities

1. Insert the Concentrix-2800 distribution tape in the Exabyte drive.
2. If the system is running in multi-user Concentrix-2800, halt the system in console mode with the shutdown command:

   ```
   # /etc/shutdown -H +5 fix root disk RETURN
   ```

   If the system is running in single-user Concentrix-2800, halt the system in console mode with the reboot command:

   ```
   # /etc/reboot -H RETURN
   ```

   If the system is powered down or the operating system is not running, power the system on (or off and on) using the power switch on the front panel. Turn the power switch to the ON (12:00) position, and press the RESET button. This stops the system in console mode.

3. Type o, B RETURN (as shown below) at the console prompts to start the system and stop in VMEbus booter mode (you can type ? for help at any of the prompts in the steps below):

   ```
   cons>Options = Enter RETURN to keep, SPACE to clear, or new op-
   tions.
   New Options? B RETURN
   ```

4. Select the tape drive as the device containing the booter program:

   ```
   cons>device boot from
   (d)isk (t)ape = t
   ```

   Note: You only need to type d and t. The system fills in the rest of the input.

5. Select the default controller for the tape drive by pressing RETURN:

   ```
   controller (0-7) [return for default] = RETURN
   ```

6. The system loads the booter. After the message announcing the Alliant VMEbus booter appears, type %r RETURN to to direct the booter to copy files onto the RAM disk, as shown below:

   ```
   Alliant VMEbus Booter [made: Wed May 9 19:32:23 EDT 1990]
   Type %h for help

   boot from:  xdv(0,0,p)dunix: %r RETURN
   ```

7. Direct the booter to load the RAM disk from the second file on the tape, as shown below:

   ```
   load ramdisk from:  xdv(0,0,p)ramdisk.init: hs(2) RETURN
   ```

8. Direct the booter to load Diagnostix-2800 from the next file on the tape, as shown below:

   ```
   boot from:  hs(2)dunix: hs(*) RETURN
   ```

9. Direct the booter to load the dddt program from the following file, at the next question, as shown below:

```
boot from:  hs(*)dddt:  hs(*) RETURN
```

10. When the Diagnostix-2800 Install Menu appears, select the desired option and press RETURN. To use the Diagnostix-2800 RAM disk utilities, choose option X:

**Diagnostix-2800 Install Menu**

```
                    Diagnostix Install Menu

      (Please Choose one of the following options and hit <return>.)

Option  Action                                              Default Command
------- ------------------------------------------------- ------------------
   I    Create root file system and Install Diagnostix    (/bin/iu.dx.sh i)
        (includes loading hard disk booter).
   U    Upgrade Diagnostix operating system.              (/bin/iu.dx.sh u)
   S    Boot to Diagnostix.                               (/etc/reboot -S)
   H    Boot to cons> prompt.                             (/etc/reboot -H)
   X    Exit to Diagnostix prompt.

Option? > X RETURN
```

The following message appears, followed by the Diagnostix Ramdisk prompt.

```
Exiting Dunix Install Utilities Menu

Diagnostix Ramdisk#
```

To exit from the Diagnostix-2800 RAM disk install utilities, type CTRL D, then one of the following options when the menu appears: S, H.

## D.3 DEVICES

The disk maintenance RAM disks support the following devices:

| | |
|---|---|
| /dev/console | console |
| /dev/dbus | diagnostic bus |
| /dev/kmem | kernel memory |
| /dev/mem | main memory |
| /dev/null | null device |
| /dev/rramdisk | physical memory disk |
| /dev/rscsh | SCSI Exabyte controller (rewind) |
| /dev/nrcsh | SCSI Exabyte controller (no rewind) |
| /dev/rxtv00h | VMEbus tape controller |
| /dev/tty | process's terminal |
| /dev/xdv00 | disk on controller 0, unit 0 |
| /dev/xdv01 | disk on controller 0, unit 1 |
| /dev/xdv10 | disk on controller 1, unit 0 |
| /dev/xdv11 | disk on controller 1, unit 1 |

## D.4 MOUNTING THE CONCENTRIX-2800 ROOT DISK

If you are using the Concentrix-2800 RAM disk, mount the root disk using one of the two available mount points, /mnt or /tmp, for example:

```
# /etc/mount /dev/xdv00r /tmp
```

## D.5 MOUNTING THE DIAGNOSTIX-2800 ROOT DISK

If you are using the Diagnostix-2800 RAM disk, mount the root disk using one of the two available mount points, /mnt or /tmp, for example:

```
# /etc/mount /dev/xdv00p /mnt
```

# INDEX


>
  diskf, A-34
?
  diskf, A-34
  lpc, A-114
  timedc, A-299
4.2BSD
  file system, A-124
8mm cartridge device, 3-21
8mm cartridge
  initializing, 3-26
  loading, 3-25
  unloading, 3-26

## A

-A
  tunefs, A-306
-a
  arp, A-5
  export client, A-52
  kgprof, A-110
  mount, A-151
  quot, A-185
  quotacheck, A-186
  quotaon/off, A-187
  repquota, A-197
  sa, A-239
  trpt, A-302
  trsp, A-304
  tunefs, A-306
  umount, A-308
abort
  lpc, A-114
  print job, 6-6
ac
  administration command, 2-2, 2-11, A-2
access
  automatic to system, 7-13
  export client, A-53
  mode for device, 3-4
account
  equivalent, 7-13
  trusted, 7-13
accounting, 2-10
  connect-time, 2-11
  custom resource, 2-12
  disable, A-3
  enable, A-3
  filter, 6-14
  login, A-2
  printer, 6-15
  printer/plotter, A-178
  resource, 2-10
  system files, A-239
  system, A-3
acct file, 1-10, 2-1, 2-10
accton
  administration command, 2-2, 2-10, A-3
acdirmax
  mount, A-153
acdirmin
  mount, A-153
acregmax
  mount, A-153
acregmin
  mount, A-153
actimeo
  mount, A-153
active
  system, 5-3
aculog file, 1-10
adb assembler, B-60
adb
  Concentrix command, 5-2
add
  diskless boot parameter, A-60
  diskless client, A-264
  NIS to network, B-20
  parallel printer to system, 6-11
  remote printer, 6-13
  remote system, B-18
  restore, A-200
  serial printer, 6-12
  user partition, B-55
  user to system, B-13
address
  Ethernet device, 7-8
  Internet, 7-10
  network, 7-10
  range
    input/output, 3-7
  resolution
    display and control, A-5
  switch, 1-7
Add_ethernet
  administration procedure, 7-3, B-2
Add_printer
  administration procedure, 6-1, B-6
Add_quotas
  administration procedure, 2-2, B-8
Add_terminal
  administration procedure, 3-2, B-10
Add_user
  administration procedure, 2-2, B-13
Add_uucp
  administration procedure, 7-3, B-17
Add_yp
  administration procedure, 7-3, B-20
adm directory, 1-9
administer
  printers, 6-5
administration command, A-133
  ac, 2-2, 2-11, A-2
  accton, 2-2, 2-10, A-3
  arcv, A-4
  arp, A-5
  automount, 7-2, A-7
  batchd, 6-1, 6-17, A-12
  biod, A-14
  bootparamd, A-15
  catman, A-16
  chown, 2-2, 2-5, A-18
  clcr, 5-1, A-19
  clinfo, 5-1, A-21
  clri, A-22
  clrm, 5-1, A-23
  compile, A-24
  comsat, 2-2, 7-2, A-25
  create_root, A-26
  cron, 5-1, 5-30, A-27
  dfmon, A-29
  diskf, 3-1, 3-12, A-30
  dlabel, 3-1, 3-14, 3-18, A-36, B-51
  dmesg, 5-1, A-40
  dump, 4-1, 4-7, A-41
  dumpdefs, A-45
  dumpfs, 4-1, 4-21
  dumpterm, A-49
  edquota, 2-2, A-50
  exportfs, A-52
  extendswap, 3-1, 3-19, A-55
  extracting, A-54
  fastboot, 1-2, A-56
  fasthalt, 1-2, A-58
  fingerd, A-59
  fix_bootparams, A-60
  fsck, 1-23, 4-1, 4-22, A-61
  fsirand, A-81
  ftpd, 7-2, A-82
  gettable, A-85
  getty, 2-2, 3-1, A-86
  gettytab, 3-42

halt, 1-2, A-88
htable, A-89
ifconfig, 7-2, 7-11, A-90
ifstats, 7-2, A-93
implog, A-95
implogd, A-96
inetd, A-97
init, 1-2, 3-1, A-100
INSTALL, A-102
keyenvoy, A-106
keyserv, 7-2, A-107
kgmon, A-108
kgprof, A-110
lockd, A-113
lpc, 6-1, 6-5, A-114
lpd, 6-1, A-116
mailq, 2-2, A-119
mailstats, 7-2, A-120
makedbm, 7-2, A-121
makefs, 3-1, 4-1, 4-5, 4-21, 4-23, A-124, B-58
MAKEHOSTS, A-121
makekey, A-128
makeramdisk, 3-1, A-129
makewhatis, A-132
make_dirs, A-123
mconnect, 7-2
mkdisknod, 3-1, 3-4, A-134
mkfile, A-135
mkhosts, A-136
mklost+found, 4-1, A-137
mknewdevs, A-138
mknod, 3-1, 3-4, 3-34, 3-36, 4-1, 7-2, 7-9, A-142
mkpasswd, 2-2, 2-4, A-143
mkproto, A-144
mkptynod, 7-9, A-146
mkstripe, 3-1, 4-1, A-147
mkswap, A-148
mktape, 3-1, 3-4, 3-28, 3-34, A-149
mkttynod, 3-1, 3-4, 3-36, A-150
mount, 3-1, 4-1, 4-3, 4-5, 7-2, A-151
mountd, 7-2, A-155
named, A-156
ndbootd, A-161
netdae, 6-1, 6-17, A-162
netrcv, 6-1, A-163
netrecv, 6-17
netsend, A-164
newaliases, 2-2, A-165, B-49, B-70
newkey, 7-2, A-166
nfsd, A-168
nfsstat, A-169
nslookup, A-171
ntalkd, 7-2, A-176
opt_software, A-177
pac, A-178
patchunix, A-179
ping, 7-2, 7-8, 7-14, A-182
portmap, 7-2, A-183
ptynum, 7-9, A-184
quot, 4-1, A-185
quotacheck, 2-2, 4-1, A-186
quotaon, 4-1, A-187
quotaon/off, 2-2
rc, A-189
rdate, 7-2, A-191
rdump, 7-2, 7-8, A-192
reboot, 1-2, A-193
renice, A-195
reportcfg, A-196
repquota, 2-14, 4-1, A-197
restore, 4-1, 4-7, A-198
rexd, A-203
rexecd, 7-2, A-204
rlogind, 7-2, A-206
rmail, A-209
rmt, 7-2, A-210
route, 7-2, A-212
routed, 7-2, A-214
rpcinfo, 7-8, A-217
rquotad, A-215
rrestore, 7-2, 7-8, A-216
rshd, 7-2, A-221
rstatd, A-223
rtnamed, A-224
rtncrtbl, A-226
rtndeltbl, A-230
rtndumptbl, A-225
rtnenter, A-230
rtnlookup, A-232
rtnpurge, A-233
rtnremove, A-234
rusersd, A-235
rwalld, 7-2, A-236
rwhod, 7-2, A-237
sa, 2-2, 2-10, A-239
sendmail, 2-2, 7-2, A-241, B-49
setclass, 5-1, 5-13, A-246
setclassdef, 5-1, A-248
setcomplex, 5-17, A-250
setcres, A-252
setdefclass, 5-13, A-254
setmaxshare, A-256
setpn, A-258
setsched, 5-1, 5-22, A-259
settype, A-263
setup_client, A-264
setup_exec, A-266
showdefclass, A-267
showmount, 7-2, A-269
showsched, 5-1, 5-15, 5-22, A-270
shutdown, 1-2, A-271, B-45
software-install, A-274
spray, A-277
sprayd, A-278
srboot, A-275
srcreate, A-281
srdelete, A-282
srdisplay, A-283
statd, A-284
sticky, 1-2
sunboot, A-286
sync, 4-1, 4-22, A-290
syslog, 7-2
syslogd, 2-2, A-292
telnetd, 7-2, A-294
tftpd, 7-2, A-295
timed, A-297
timedc, A-299
timezone, A-301
tnamed, A-296
trpt, A-302
trsp, A-304
tunefs, 4-1, 4-21, 4-23, A-306
umount, 4-1, 4-6, A-308
unmkstripe, 4-1, 4-6, A-310
update, 4-1, 4-22, A-311
uucico, 7-2, 7-17, A-312
uuclean, 7-2, 7-17, A-315
uucp, 7-17
uucpd, 7-2, 7-17, A-316
uupoll, 7-2, 7-17, A-317
uusnap, 7-2, 7-17, A-318
uux, 7-17
uuxqt, 7-2, 7-17, A-319
uv, A-320
verify_tapevol_arch, A-321
vipw, 2-2, 2-4, A-322
vmlock, 5-1, A-323
xdrtoc, A-324
xdvstat, A-325
ypbind, 7-2, 7-19, A-326
ypinit, 7-2, A-330
ypmake, 7-2, A-331
yppasswd, 7-2, A-333
yppoll, 7-2, A-334
yppull, A-335
yppush, 7-2, A-336
ypserv, 7-2, 7-19, A-338
ypset, A-340
ypupdated, 7-3, A-341
ypwhich, 7-2, A-342
ypxfr, 7-3, A-343
zic, A-345
administration procedure
Add_ethernet, 7-3, B-2
Add_printer, 6-1, B-6
Add_quotas, 2-2, B-8
Add_terminal, 3-2, B-10
Add_user, 2-2, B-13
Add_uucp, 7-3, B-17
Add_yp, 7-3, B-20
Boot_system, 1-2, B-23
Broadcast, 2-2, B-34
Build_kernel, 5-1, B-36
Call_cs, 1-2, B-37
Connect_diag, 1-2, B-38
Disable_terminal, B-39
Extend_swap, 3-2, B-40

Format_disk, 3-2, 4-2, B-42
Halt_system, 1-2, B-45
Install_batch, 6-1, B-46
Install_mail, 2-2, B-49
Label_disk, 3-2, 4-2, B-51
Make_file_system, 4-2, B-58
Patch_kernel, 5-1, B-60
Recover_system, 1-2, B-66
Remove_user, 2-2, B-70
Shutdown_system, 1-2, B-73
af printcap entry, 6-14
aF
diskf, A-31
afio
Concentrix command, 4-2
ALIASES
ypmake, A-331
aliases, 2-8
format of file, 1-11, 2-1, B-16, B-50
rebuild mail data base, A-165
ypmake, A-331, B-20
aliases.dir file, 2-1
aliases.pag file, 2-1
all
setclassdef, A-248
Alliant ddt 2800 mode, 1-21
Boot_system, B-31
Alliant ddt mode, 1-17, B-28
Alliant ddt, 1-13
alliant directory, 1-9
Alliant
customer service, 1-39, B-66
report problem to, B-37
alternate superblock
file system, 4-26
tunefs, A-306
am
diskf, A-31
anon
export client, A-52
anritsu tape drive, 3-27
any default class
Scalar/C or no concurrency, 5-13
any
execution mode, 5-12
setclass, A-246
setdefclass, A-254
setsched, A-261
showdefclass, A-267
ap
diskf, A-31
Apollo window
terminal, 3-43
architecture
diskless install, A-102
archive files
convert, A-4
archiver
tape, 4-15
arcv
administration command, A-4
arp
administration command, A-5
ifconfig, A-90
special file, 7-4
-arp
ifconfig, A-90
assign
computing resource type, A-252
asynchronous block I/O
daemon, A-14
attach processors, 1-30
attached
CE, 5-13
au
diskf, A-31
aus
timezone, A-301
AUTH_DES
ypupdated, A-341
AUTH_UNIX
ypupdated, A-341
autoboot
front panel, 1-3
automatic system access, 7-13
automount
administration command, 7-2, A-7
aZ
diskf, A-31


## B


B
diskf, A-31
b
diskf, A-31
dump, A-41
mknod, A-142
restore, A-199
-B
fastboot, A-56
fsck, A-61
makefs, A-124
reboot, A-193
shutdown, A-271
-b
dlabel, A-36
fastboot, A-56
fsck, A-61
kgmod, A-108
kgprof, A-110
makedbm, A-121
makefs, A-124
named, A-156
reboot, A-193
rpcinfo, A-217
sa, A-239
-ba
sendmail, A-241
backup
incremental file system, A-41
bad block
diskf, A-31
bad sector
SCSI disk, 3-23
basic operations, 1-1
Batch
Device Descriptions, 6-17, B-46
Mappings, 6-17, B-47
Parameters, B-46
Queue Descriptions, 6-17, B-47
batch
configuration file, 6-17
daemon, A-12
delete request, 6-18
directory, 6-1
install, B-46
log file, 1-10
logical device, 6-16
network receive daemon, A-163
network send daemon, A-164
network server daemon, A-162
queue, 6-1, 6-16
setup, 6-17
status, 6-18
submit request, 6-17
system
queue, A-12
batchd
administration command, 6-1, 6-17, A-12
-bd
sendmail, A-241
bfio
Concentrix command, 4-2, 4-14
bg
mount, A-152
-bi
sendmail, A-241
biff
Concentrix command, 2-3, 7-3
Internet service, 7-6
server, A-25
/bin directory, 1-9
bind
NIS, A-326
biod
administration command, A-14
block device
RAM disk, A-129
block
blocks per transfer, 4-23
data
file system, 4-24

device mode, A-142
file system, 4-24
indirect pointer
file system, 4-23
logical number, A-62
size, A-124, 4-23
blocking factor
dump, A-41
restore, A-199
-bm
sendmail, A-241
boot flags, 1-23
srboot, A-275
boot from RAM disk, 1-37
boot loop
console, 1-18
boot script, 1-25
quotas, 1-33
boot
diskless server, A-161
flag, B-25
parameter server, A-15
partition, 3-14
Sun system, A-286
boot/device
front panel, 1-3
booter program, 1-23
booter, A-275
Boot_system, B-25
Concentrix-2800, 1-13
Diagnostix-2800, 1-13
reboot, A-196
bootparamd
administration command, A-15
Boot_system
administration procedure, 1-2, B-23
Alliant ddt 2800 mode, B-31
alliant ddt mode, B-28
Concentrix-2800 mode, B-32
console mode, B-26
diagnostic monitor mode, B-30
Diagnostix-2800 mode, B-28
FX/2800 mode, B-31
FX/2800 startup mode, B-29
single-user mode, B-32
VMEbus booter mode, B-27
bourne
shell, B-15
-bp
sendmail, A-241
brand
tape drive, 3-5
breaktop kernel variable, B-62
Broadcast
administration procedure, 2-2, B-34
broadcast
ifconfig, A-90
messages to users, B-34
-bs
sendmail, A-241
BSD
file system, 4-4
-bt
sendmail, A-241
buffer
for Ethernet performance, 7-15
build
NIS, A-330
special file, A-142
Build_kernel
administration procedure, 5-1, B-36
-bv
sendmail, A-241
bypass
file system checks, 1-14
byte interface
Concentrix to Diagnostix, 3-6
bytes
per i-node
file system, 4-23
-bz
sendmail, A-241


## C


C
diskf, A-31
rmt, A-210
c
diskf, A-34
dump, A-41
mknod, A-142
-C
sendmail, A-241
-c
fastboot, A-56
implog, A-95
kgmod, A-108
kgprof, A-110
makefs, A-124
nfsstat, A-169
pac, A-178
quot, A-185
reboot, A-193
sa, A-239
setcomplex, A-250, A-250
spray, A-277
cabling, 3-38
cache module
hardware identification, 1-7
cache
IP, 1-5
cacstack kernel variable, B-62
cacstacksegs kernel variable, B-62
cacstacksize kernel variable, B-62
call remote file, 7-18
Call_cs
administration procedure, 1-2, B-37
cartridge tape system
SCSI, 3-24
cartridge
VMEbus, 3-6
catman
administration command, A-16
cce default class
computational complex, 5-13
cce
execution mode, 5-12
setclass, A-246
setdefclass, A-254
setsched, A-261
showdefclass, A-267
cd
restore, A-200
CE
default scheduling vector, 5-14
dynamically detached, 5-17
hardware identification, 1-4
scheduling, 5-11
setsched, A-259
statically detach, 5-13
cf
print filter, 6-14
change
computing resource type, A-263
default devices, 3-36
execution mode, 5-12
process priority, A-195
character device mode, A-142
characteristics
terminal, B-10
chassis ground, 3-38
check
file system, 1-23, 4-22
chgrp
Concentrix command, 2-3, 2-5, 4-2
chkey
Concentrix command, 2-3
chmod
Concentrix command, 2-3, 2-5, 4-2, 7-3
chown
administration command, 2-2, 2-5, A-18
cifplot filter, 6-14
cipher tape drive, 3-27
CL
default scheduling vector, 5-14
scheduling, 5-11

setcomplex, A-250
setsched, A-259
class
process, 5-13
scheduling, 5-13
set default scheduling, A-254
set process, A-246
clcr command, 1-30
clcr
administration command, 5-1, A-19
clean
directory spool, 7-17
lpc, A-114
uucp spool directory, A-315
cleaning up /tmp, 4-19
clear
i-node, A-22
client
diskless install, A-102
NIS, 7-20, B-21
clinfo
administration command, 5-1, A-21
clock
daemon, 5-30, A-27
hardware identification, 1-7
clockdiff
timedc, A-299
clri
administration command, A-22
clrm command, 1-30
clrm
administration command, 5-1, A-23
cluster
computational, 5-16
configuration
startup file, 5-21
startup, 5-19
create, A-19
different sizes, 5-26
dynamic reconfiguration, 5-20
dynamics, 5-17
kill, 5-19
print information, A-21
processors, 5-11
remove, A-23
scheduling, 5-11
CM
configuration, 1-7
coalesce
computational cluster, 5-17
computational complex, A-250
coaxial cable, 7-5
command
script to start system, A-189
unix to unix execution, 7-17
uux can execute, 7-18
communication
user, 2-7
compile
administration command, A-24
terminfo database, A-24
composition of cluster
print information, A-21
computational cluster, 5-16
create, A-19
delete, A-23
print information, A-21
computational complex
different number of CEs, 5-17
ducomplex, 5-17
set threshold for, A-250
set threshold, A-250
computational configurations, 1-6
computational element, 1-4
attached, 5-17
cache, 1-5
dynamically detached, 5-17
scheduling, 5-11
statically detach, 5-13
computational resources
scheduling, 5-11, 5-13
computing resource
set, A-252
comsat
administration command, 2-2, 7-2, A-25
Concentrix command
adb, 5-2
afio, 4-2
bfio, 4-2, 4-14
biff, 2-3, 7-3
chgrp, 2-3, 2-5, 4-2
chkey, 2-3
chmod, 2-3, 2-5, 4-2, 7-3
dd, 3-31
df, 4-2, 4-3, 5-2, 5-9
domainname, 7-3, 7-19
du, 5-2, 5-9
execute, 5-2
find, 5-2, 5-10
finger, 2-6
ftp, 7-3
grep, 5-2, 5-6
hostname, 7-3
kconfig, 5-2
keylogin, 2-3
kill, 1-39, 5-2, 5-8, 7-3
login, 2-3, 2-3
logout, 2-3, 2-3
lpq, 6-2, 6-4, 6-5
lpr, 6-2, 6-4
lprm, 6-2, 6-4, 6-5
ls, 1-12, 4-2, 4-3, 5-10
mail, 2-3, 2-7
mkdir, 2-3, 4-2, 4-5
mknod, 3-4
mon, 5-2, 5-4
motd, 2-7
msgs, 2-3, 2-7
netstat, 7-3, 7-8, 7-14
passwd, 2-3
pr, 6-2
print, 6-2, 6-4
ps, 1-39, 5-2, 5-6
qdel, 6-2, 6-18
qstat, 6-2, 6-18
qsub, 6-2, 6-17
quota, 2-3, 4-2
rcp, 7-3
rlogin, 2-3, 2-3, 7-3
rmail, 7-3
rsh, 7-3
rt, 3-31
ruptime, 7-3
rwall, 7-3
rwho, 7-3
su, 1-8, 2-3, 2-3
syscfg, 3-10, 5-2, 5-3
talk, 7-3
tar, 4-2, 4-15
telnet, 2-3, 2-3, 7-3
tftp, 7-3
tip, 7-3
umask, 2-6
uucp, 7-3
uulog, 7-3
uuname, 7-3
uuq, 7-3
uusend, 7-3
uux, 7-3
w, 5-2, 5-3
wall, 2-3, 2-7
xget, 2-3
xproc, 5-2
xsend, 2-3
ypcat, 7-3, 7-19
Concentrix
boot, A-196
multi-user, 1-8
patch image, A-179
single-user, 1-8
Concentrix-2800 mode, 1-22
Boot_system, B-32
Concentrix-2800 root partition, 1-25
Concentrix-2800, 1-13
configuration file, 5-18
RAM disk utilities, D-1
root partition
restoring, 4-18
single-user mode, 1-21
start up, 1-23
concurrency
clusters, 5-26
execution mode, 5-12
concurrent job
set max processors, A-258
concurrent jobs
high use, 5-23
concurrent
job
setcomplex, A-250

conf.c file, 3-1
config file
  srboot, A-275
configuration capabilities, 1-6
configuration file
  cluster specifications, 5-19
configuration options
  sendmail, A-242
configuration
  clusters, 5-16
  hardware, 1-4, 5-11
  inetd file, A-98
  network, 7-11
  report software/hardware, A-196
  software, 1-8
  system, 5-3
configure
  Apollo window as terminal, 3-43
  batch system, 6-17
  clusters
    rc.cluster, 5-21
  network interface parameter, A-90
  special tape drive, 3-34
  tape drive, 3-28
connect
  to Diagnostix-2800, B-38
  to SMTP mail server socket, A-133
connect-time
  accounting, 2-11
connection
  test on network, 7-14
connection-oriented
  inetd, A-97
connectivity
  fsck, A-66
Connect_diag
  administration procedure, 1-2, B-38
consistency check
  file system, 4-22, A-61
console boot loop, 1-18
console mode, 1-14
  Boot_system, B-26
  help command, 1-15
  options command, 1-15
  options, 1-13
  parameters command, 1-15
  unix command, 1-15
console
  configuration, 5-1
  connect to Diagnostix-2800, B-38
  device, 3-6
  directory, 1-9
  duart port, 3-6
  file, 1-1
  halt system in, A-88
  mode, 1-23, B-45
  port name, 3-35
  port, 3-6
  printer connection, 3-39
  terminal, 1-3
console.config file, 5-1
console.config
  startup, 5-18
/console/console.config
  startup, 5-18
console_loop
  srboot, A-275
construct
  file system, A-124
control terminal
  rlogind, A-206
control
  line printer, 6-1
controller
  CMC Ethernet, 3-6
  disk, 3-11, 3-12
  Excelean Ethernet, 3-6
  making device, 3-5
convert
  archive files, A-4
  distribution table of contents, A-324
  NIC host table, A-89
cooked device mode, A-142
copy
  unix to unix, 7-17
core dump, 5-10
corrupted root disk, D-1
cpu
  top users, 5-4
-Cr
  dlabel, A-36
create
  cluster, 5-20, A-19
  file system, B-58
  logical name table, A-226
  new key
    publickey data base, A-166
  NFS swap file, A-135
  on-line documentation, A-16
  RAM disk image, A-129
  shared memory region, A-281
create_root
  administration command, A-26
cron
  administration command, 5-1, 5-30, A-27
crontab file, 1-11, 2-10, 5-30
crontab
  file, 5-1
crossbar interconnect, 1-7
cs directory, 1-9
-Cs
  dlabel, A-36
csh shell, B-15
.cshrc file, 2-6
ct
  timezone, A-301
CTRL/C
  hung process, 5-8
CTRL/Q
  hung process, 5-8
CTRL/Z
  hung process, 5-8
CTS function
  cabling, 3-38
current users on system, 5-3
custom
  kernel, B-36
  resource accounting, 2-12
customer service, 1-39, B-66
  report problem to, B-37
customunix
  custom kernel file, B-36
cylinder
  cylinders per group, A-124
    file system, 4-23
  disk, 3-12, 3-14
  group information, A-45
  group
    tunefs, A-306

## D

D
  diskf, A-31
d
  dump, A-41
-D
  fastboot, A-56
  implog, A-95
  reboot, A-193
  sa, A-239
  shutdown, A-271
-d
  ac, A-2
  arp, A-5
  bootparamd, A-15
  cron, A-27
  fastboot, A-56
  inetd, A-97
  makedbm, A-121
  named, A-156
  ndbootd, A-161
  netdae, A-162
  nfsstat, A-169
  reboot, A-193
  routed, A-214
  rpcinfo, A-217
  sa, A-239
  sendmail, A-241
  shutdown, A-271
  spray, A-277
  syslogd, A-292
  uucico, A-312
  yppush, A-336
  ypset, A-340
  ypwhich, A-342
daemon
  automount
    NFS filesytem, A-7
  batch network receive, A-163
  batch network send, A-164

batch network server, A-162
batch, A-12
clock, 5-30, A-27
Internet services, A-97
kernel statistics server, A-223
key server, A-107
line printer, A-116
log system messages, A-292
logical name, A-224
name server, A-156
network lock, A-113
network routing, A-214
network status monitor, A-284
network username server, A-235
NFS mount request server, A-155
NFS, A-14, A-168
NIS binder process, A-326
NIS password file server, A-333
remote execution server, A-204
remote login server, A-206
remote quota server, A-215
remote shell server, A-221
RPC remote execution server, A-203
rwall server, A-236
server for changing NIS info, A-341
spray server, A-278
talk program, A-176
telnet protocol server, A-294
time server, A-297
trivial name server, A-296
UUCP, A-316
daily script, 2-10, 5-30, 7-18
DARPA
Internet file transfer protocol server, A-82
port to RPC program number mapper, 7-6, A-183
TELNET protocol server, A-294
trivial file transfer protocol server, A-295
trivial name server, A-296
data base
NIS, 7-19
data switch, 1-7
data
block
file system, 4-24
communication equipment (DCE), 3-38
terminating equipment (DTE), 3-38
datagram
user, 7-5
datagram-oriented
inetd, A-97
date
cron execution, A-27
set for system
from remote host, A-191
set, 1-38
dbm
NIS, A-121
DCD function
cabling, 3-38
DCD
switch, 3-40
dce default class
reserved, 5-13
DCE mode
cabling, 3-38
dce
setclass, A-246
setdefclass, A-254
setsched, A-261
showdefclass, A-267
dd
Concentrix command, 3-31
/dddt, 1-16
ddt 2800
startup, B-31
ddt mode
Alliant, 1-17
options, 1-17
ddt
srboot, A-275
/ddt
booter, 1-20
ddtbrkflag kernel variable, 1-40, 1-40
debug
ifconfig, A-90
-debug
ifconfig, A-91
debugging
Diagnostix-2800, B-27
network, A-182
dec
execution mode, 5-12
DECnet driver
ge , 3-6
dedicate resource, 5-28
default devices
changing, 3-36
create, A-138
pseudo terminals, 7-9
default partitions
file system, 4-4
default
display scheduling classes, A-267
printer, 6-10
scheduling vector, 5-14
set scheduling class, A-254
special file, 3-10
delete
batch request, 6-18
cluster, A-23
diskless boot parameter, A-60
logical name table, A-230
restore, A-200
shared memory region, A-282
density
dump, 4-10, A-41
tape, 3-5
desfree kernel variable, B-62
detach processors, 1-30
detached jobs
high use, 5-25
detached
CE, 5-13, 5-17
execution mode, 5-12
/dev directory, 1-9, 3-1, 3-2, 6-1
permissions, 3-18
device, 3-1
access mode, 3-4
changing defaults, 3-36
create defaults, A-138
Diagnostix-2800, 3-8
disk, 3-11
Ethernet address, 7-8
full, 3-6
list of supported, 3-5
logical batch, 6-16
major number, 3-2
minor number, 3-2
name, 3-4
Ethernet, 7-8
null, 3-6
RAM disks, D-9
SCSI disk, 3-6
srboot, A-275
tape, 3-27
terminal, 3-35
df
Concentrix command, 4-2, 4-3, 5-2, 5-9
print filter, 6-14
dfmon
administration command, A-29
Diagnostic monitor mode
Boot_system, B-30
diagnostic monitor, 1-18
diagnostic
partition, 3-14
diagnostics
restore, A-201
rexecd, A-205
rlogind, A-208
rshd, A-222
sendmail, A-244
Diagnostix
boot, A-196
fastboot, A-56
Diagnostix-2800 booter
Boot_system, B-25
Diagnostix-2800 devices
table, 3-9
Diagnostix-2800 mode, 1-17
Boot_system, B-28
Diagnostix-2800 root partition, 1-24

Diagnostix-2800, 1-13, 1-31
  configuration file, 5-18
  devices, 3-8
  install utilities
    RAM disk, D-3
  loading, 1-16
  Maintenance Utilities, D-2
  memory map, 3-48
  RAM disk image, A-129
  root partition
    restoring, 4-16
dict directory, 1-9
DIR
  ypmake, A-331
directory
  clean up spool, 7-17
  list contents, 4-3
  list of system, 1-9
  list, 1-12
  make, A-123
  move with bfio, 4-14
  move with tar, 4-15
  pathname, 4-3
  save with bfio, 4-14
  save with tar, 4-15
  sticky bit, A-285
disable
  disk quotas, A-187
  lpc, A-114
  print queueing, 6-6
  swapping/paging, A-323
  system accounting, A-3
  terminal, B-39
Disable_terminal
  administration procedure, B-39
disassemble
  disk stripe, A-310
disconnect
  from Diagnostix-2800, B-38
disk quotas
  file system, 2-13
disk stripes
  file system, 4-4
disk
  controller, 3-11
  devices, 3-11
  disassemble stripe, A-310
  format, B-42
  geometry, 3-12
  hardware flaw map, 3-16
  initialization
    examples, A-34
  internal flaw map, 3-16
  label, A-36, B-42, B-51
  labels, 3-13
    SCSI, 3-22
  layout, 3-17
  maintenance, D-1
  make special file, A-134
  name, 3-11
  partition, 3-13
  protection, 3-18
  quotas
    check, A-186
  report free space, 5-9
  SCSI, 3-6
  storage layout, 3-12
  summary and usage
    dynamic, A-29
  supported controllers, 3-11
  supported drives, 3-11
  usage, 5-9, A-196
    summary for file system, A-197
  user space, 3-17
  utility program, A-30
  VMEbus, 3-6, B-42
  xdv
    display statistics, A-325
diskf
  administration command, 3-1, 3-12, A-30
DISKF_SCRATCH_DISK, 3-18
diskless client
  create root, A-26
diskless
  add client, A-264
  boot block server, A-161
  boot parameter server, A-15
  boot parameter, A-60
  extract distribution, A-54
  install architecture and clients, A-102
  install client software, A-266
  remove client, A-264
  select software from distribution, A-177
  verify mounted distribution tape, A-321
disktab
  disk label, A-36
  enhanced version, B-52
  file, 1-11, 3-1, B-51
display
  cluster configuration, 5-20
  default scheduling classes, A-267
  disk usage
    dynamic, A-29
  file system information, 4-21
  kernel version information, A-320
  process times, 5-7
  processors activity, 5-7
  resource scheduling vector, A-270
  shared memory region, A-283
  system memory usage, 5-5
  system status, 5-4
  xdv disk statistics, A-325
/dist directory, 1-9
distribution
  diskless, A-54
ditroff filter, 6-14
dlabel
  administration command, 3-1, 3-14, 3-18, A-36, B-51
dmesg
  administration command, 5-1, A-40
DNX-2800
  DECnet driver, 3-6
documentation
  on-line, A-16
DOM
  ypmake, A-331
domain
  name server, 7-6, A-156
  name, 7-19
domainname
  Concentrix command, 7-3, 7-19
down
  ifconfig, A-91
  lpc, A-114
-Dr
  dlabel, A-36
driver
  disk, 3-12
drop cable, 7-5
-Ds
  dlabel, A-36
DSR function
  cabling, 3-38
DSR
  switch, 3-40
dstaddr
  ifconfig, A-91
DTE mode
  cabling, 3-38
DTR function
  cabling, 3-38
DTR
  switch, 3-40
du
  Concentrix command, 5-2, 5-9
ducomplex
  computational complex, 5-17
duart port, 3-6, 3-37
dump
  across network, A-192
  administration command, 4-1, 4-7, A-41
  density, 4-10
  diskless, A-102
  file system across network, 7-8
  file system information, 4-21, A-45
  file system, 4-7
  level, 4-7
  logical name table, A-225
  media, 4-10
  procedure, 4-8
  restore across network, 7-8
  script, 4-9
  size, 4-10
dumpdates

dump, A-42
dumpfs
administration command, 4-1, 4-21, A-45
dumpterm
administration command, A-49
dunix prompt, 1-17
dunix shell, B-28
dunix
file, 1-1
/dunix, 1-16
duu directory, 1-9

## E

e
dump, A-41
restore, A-199
-E
kgprof, A-110
-e
arp, A-5
kgprof, A-110
tunefs, A-306
echo
icmp echo packet, A-182
icmp packet, 7-8
edit
password file, A-322
time limits
file system, A-50
user quotas, A-50
edquota
administration command, 2-2, A-50
ee
timezone, A-301
elog file, 5-1
en
CMC Ethernet controller, 3-6
enable
disk quotas, A-187
print queueing, 6-6, A-115
swapping/paging, A-323
system accounting, A-3
terminal, B-39
enbproflag kernel variable, A-108, B-62
encrypt
mail, 2-8
encryption
generate key, A-128
enter
logical name in table, A-230
environment variable
PRINTER, 6-4
equivalent systems and accounts, 7-13
ERRLOG file, 1-10
error
fsck, A-67
peripheral, 1-41
print messages, 6-7
et
timezone, A-301
/etc directory, 1-9
/etc/rc.cluster, 5-21
Ethernet controller, 3-6
ex, 3-6
Ethernet
add system to, B-2
address, 7-9
board statistics, A-93
controller name, 7-8
controller, 3-6
driver placement, 7-15
performance, 7-15
protocols supported, 7-5
translation in arp, A-5
ethers
Add_ethernet, B-2
file, 1-11, 7-1, 7-11
ypmake, A-331, B-20
ex
Excelan Ethernet controller, 3-6
Exabyte
8mm cartridge device, 3-21
initializing, 3-26
loading, 3-25
SCSI helical scan cartridge, 3-24
unloading, 3-26
examine
network, 7-14
print queue, 6-4
tape, 3-31
example
disk initialization, A-34
dynamic cluster reconfiguration, 5-20
exec routine
scheduling class, 5-13
exec
diskless, A-102
Internet service, 7-6
executable file, 5-10
execute command
scheduling class, 5-13
execute
Concentrix command, 5-2
force execution mode, 5-12
execution mode
change, 5-12
scheduling, 5-12
execution
remote program, 7-7
remote, 7-6
unix to unix command, 7-17
uucp file interpreter, A-319
exercise
disk, A-30
exit
lpc, A-115
explode
computational cluster, 5-17
computational complex, A-250
export
NFS client, A-52
exportfs
administration command, A-52
exports
file, 7-1, A-52
extend
swapping/paging area, B-40
extendswap
administration command, 3-1, 3-19, A-55
Extend_swap
administration procedure, 3-2, B-40
extract
restore, A-200
extracting
administration command, A-54
ex_n_pages kernel variable, B-62

## F

F
diskf, A-31
f
diskf, A-31
dump, A-42
restore, A-199
-F
kgprof, A-110
makefs, A-124
sendmail, A-241
-f
arp, A-5
chown, A-18
dlabel, A-36
implog, A-95
kgprof, A-110
makefs, A-124
mount, A-151
quot, A-185
route, A-212
sa, A-239
sendmail, A-241
shutdown, A-271
syslogd, A-292
trpt, A-302
failure
system, 1-39
fastboot, 1-14, 1-23, A-271
administration command, 1-2, A-56
Boot_system, B-23
file, 1-1
fasthalt
administration command, 1-2, A-58
FDDI
pg device, 3-6

fg
  mount, A-152
fields
  partition table
    Label_disk, B-52
file system, 1-9
  4.3BSD, 4-4
  Add_quotas procedure, B-8
  backup, A-41
  block, 4-24
  bypass checks, 1-14
  check quotas, A-186
  check, 1-23
  configurations, 4-5
  consistency check, 4-22
  construct, A-124
  disk quotas, 2-13
  dump across network, 7-8
  dump information, 4-21, A-45
  dump, 4-7
  files and commands, 4-1
  fragment, 4-24
  free space, 4-25
  i-node, 4-24
  incremental restore, A-198
  layout, 4-23
  make and mount, B-58
  make, 4-4, 4-23
  mount, A-151
  move with dump, 4-13
  optimization, A-306
  prototype, A-144
  quotas example, 2-13
  remote dump, A-192
  remote mount, 7-6
  remote restore, A-216
  repair, A-61
  restore, 4-7, 4-11
  root, 3-17, 4-3
  rotational delay, 4-26
  striped, 4-4, B-57, B-58
  summarize ownership, A-185
  superblock, 4-22
  synchronization, 4-22
  threshold, A-306
  track-allocated, 4-4
  tune, 4-21, 4-23, A-306
  turn quotas on/off, A-187
  UNIX, 1-8
  unmount, A-308
  usr, 4-3
file
  change owner, A-18
  executable, 5-10
  limit in cylinder, 4-23
  list, 1-12, 5-10
  needing maintenance, 1-11
  object, 5-10
  remote calling information, 7-18
  special, 3-2
  system files that grow, 1-10, 5-10
  temporary, 5-10
  transfer via uucp, 7-17
  transfer, 7-6
  trivial file transfer program, 7-7
  trivial file transfer protocol server, A-295
find
  Concentrix command, 5-2, 5-10
finger
  Concentrix command, 2-6
  daemon, A-59
  Internet service, 7-6
fingerd
  administration command, A-59
fix_bootparams
  administration command, A-60
fj1000s disk drive, B-42
fj1000s
  Fujitsu disk drive, 3-11
fj2600s
  Fujitsu disk drive, 3-11
fj2f600s disk drive, B-42
fj690s disk drive, B-42
fj690s
  Fujitsu disk drive, 3-11
flags
  boot, B-25
  modes of operation, 1-13
  srboot, A-275
flaw map, 3-16, B-42
  diskf, A-34
flaw-file
  diskf, A-31
flaws
  SCSI disk, 3-23
format
  disk, A-30, B-42
Format_disk
  administration procedure, 3-2, 4-2, B-42
Fortran filter, 6-14
.forward file, 2-1
fpr
  file, 6-1
  Fortran filter, 6-14
fragment
  file system, 4-24
  size, A-124
    file system, 4-23
free block
  fsck, A-64
free space
  disk, 5-9
  file system, 4-25
  minimum
    file system, 4-23
front panel, 1-3
fsck administration command, 1-23
fsck
  administration command, 4-1, 4-22, A-61
  lost+found, A-137
fsirand
  administration command, A-81
fstab
  file, 1-1, 1-11, 3-1, 4-1
ftp
  Concentrix command, 7-3
  Internet service, 7-6
ftpd
  administration command, 7-2, A-82
ftpusers
  file, 7-1
Fujitsu
  disk drive, 3-11, B-42
  tape drive, 3-27
full device, 3-6
full
  blocks
    file system, 4-23
FX/2800 booter mode, 1-20
  Boot_system, B-31
FX/2800 booter, 1-13, A-275
FX/2800 startup mode, 1-18
  srboot, A-275
FX/2800 startup, 1-13
  diagnostix-2800 mode, 1-17
FX/2800
  computational configurations, 1-6
FX/800
  computational configuration, 1-6
FX2800 startup mode
  Boot_system, B-29

## G

-g
  lockd, A-113
  routed, A-214
  rtncrtbl, A-226
  uucico, A-312
games directory, 1-9
gateway-gateway protocol, 7-5
gateways
  file, 7-1
GCOS, B-15
ge
  DNX-2800 DECnet driver, 3-6
generate
  encryption key, A-128
  hashed host table, A-136
  hashed password table, A-143
generic
  job
    setcomplex, A-250
get
  NIS map version, A-334
gettable

administration command, A-85
getty
administration command, 2-2, 3-1, A-86
gettytab
administration command, 3-42
file, 1-1, 1-11, 2-1, 3-1, 7-1
getty command, A-86
gf
print filter, 6-14
GLIST
SCSI disk, 3-23
grep
Concentrix command, 5-2, 5-6
group
chown, A-18
cylinder
tunefs, A-306
file, 1-1, 1-11, 2-1, 2-5, 4-1, 7-1
format of file, B-16
identification number, B-15, B-16
name, B-16
ypmake, A-331, B-20
grow
system files that, 1-10
grown flaw
SCSI disk, 3-23
grpid
mount, A-151


## H


H
diskf, A-31
h
diskf, A-31, A-34
restore, A-199
-H
fastboot, A-56
fasthalt, A-58
halt, A-88
reboot, A-193
shutdown, A-271
-h
fastboot, A-56
implog, A-95
kgmod, A-108
newkey, A-166
quot, A-185
reboot, A-193
sendmail, A-242
shutdown, A-271
umount, A-308
ypset, A-340
halt
administration command, 1-2, A-88
fast, A-58
fastboot, A-56
from multi-user mode, 1-22
multi-user mode, 1-14
system in console mode, A-88, B-45
system, A-58
Halt_system
administration procedure, 1-2, B-45
handle
remote mail from uucp, A-209
hang
recover, B-66
system, 1-40
terminal, 1-39
hangup
signal to init, 3-43
hard mount, A-152
hardware driver
Ultranet, 3-6
hardware flaw map
SCSI disk, 3-23
hardware identification
cache module, 1-7
clock, 1-7
IOM, 1-5
memory module, 1-7
PM, 1-4
switch, 1-7
hardware
configuration, 1-4, 5-11
flaw map, 3-16
network problem, 7-14
report configuration, A-196
terminal configuration, 3-37
hashed
make host table, A-136
make password table, A-143
helical scan cartridge, 3-24
Helicscan tape drive and subsystem, 3-21
help command
console mode, 1-15
help
lpc, A-115
restore, A-200
timedc, A-299
high priority
scheduling, 5-27
home
directory, B-15
diskless, A-102
host
make hashed table, A-136
NIC host table, A-85
trusted, 7-13
hostname
command, B-49
Concentrix command, 7-3
system, 7-10
hosts
Add_ethernet, B-2
directory, 1-9, A-121
file, 1-11, 7-1, 7-11
ifconfig, A-90
ypmake, A-331, B-20
hosts.equiv file, 7-13
hosts.equiv
file, 1-11, 7-1
hourly
uucp script, 7-18
hpages kernel variable, 7-16
htable
administration command, A-89
hung system
reset, 1-22
.hushlogin file, 2-1, 2-7


## I


I
rmt, A-210
i
diskf, A-31, A-34
restore, A-198
-I
fsck, A-61
makefs, A-124
-i
dlabel, A-36
export client, A-52
implog, A-95
makedbm, A-121
makefs, A-124
spray, A-277
timed, A-297
ypupdated, A-341
i-node
clear, A-22
data size in fsck, A-66
file system, 4-24
fsck, A-61
links in fsck, A-65
random generation numbers, A-81
size, A-124
state in fsck, A-65
ICMP
network protocol, 7-5
icmp
echo packet, A-182
special file, 7-4
if
print filter, 6-14
special file, 7-4
ifconfig
Add_ethernet, B-2
administration command, 7-2, 7-11, A-90
ifstats
administration command, 7-2, A-93
image memory unit adapter card, 3-6
imb
byte interface to Diagnostix-2800, 3-6
IMP

log interpreter, A-95
logger process, A-96
implicit migration to IP, 5-29
implog
administration command, A-95
implogd
administration command, A-96
imuac
image memory unit adapter card, 3-6
include directory, 1-9
incremental file system restore, 4-11, A-198
indirect
data block
file system, 4-24
pointer block size, A-124
pointer block
file system, 4-23
inet
special file, 7-4
inetd
administration command, A-97
inetd.conf
file, 1-11
init
administration command, 1-2, 3-1, A-100
NIS, A-330
initialize
system, A-100
initializing
8mm cartridge, 3-26
Exabyte, 3-26
input/output address range, 3-7
input/output module
hardware identification, 1-5
install utilities
RAM disk, D-1
INSTALL
administration command, A-102
install
default devices, A-138
diskless client software, A-266
layered product, A-274
NIS, A-330
Install_batch administration procedure, B-46
Install_batch
administration procedure, 6-1
Install_mail administration procedure, 2-2
Install_mail
administration procedure, B-49
interactive processor, 1-4
cache, 1-5
scheduling, 5-11
interactive
file system check, A-63
restore, 4-12, A-200
interface network parameter, A-90
internal flaw map, 3-16
internet driver
Ultranet, 3-6
Internet
address, 7-10, B-2
Control Message Protocol, 7-5
domain name server, A-156
file transfer protocol server, A-82
name servers
query, A-171
send mail over, 7-7
send mail, A-241
services daemon, A-97
services, 7-6
to-Ethernet translation for arp, A-5
Interphase Jaguar controller, 3-21
Interphase Peregrine FDDI
pg, 3-6
intr
mount, A-152
ioconf.c file, 3-1
IOM
hardware identification, 1-5
IP only
execution mode, 5-12
IP
cache, 1-5
default scheduling vector, 5-14
network protocol, 7-5
scheduling, 5-11
setsched, A-259
software identification, 1-5
ip
special file, 7-4
IP0
default scheduling vector, 5-15
hardware identification, 1-5
IPI-2 disk drive interface, B-42
IPI-2 disk
notes
Label_disk, B-55
IPI-2 interface, 3-11
ipo default class
daemons, 5-13
ipo
execution mode, 5-12
setclass, A-246
setdefclass, A-254
setsched, A-261
showdefclass, A-267
ISR_commsize kernel variable, B-62
ISR_loadsize kernel variable, B-62


## J


-j
sa, A-239
trsp, A-304
Jaguar
SCSI interface, B-42
job
batch, 6-16, B-46
concurrent
set processors, A-258
kill, 5-8
print, 6-3


## K


-K
sa, A-239
-k
sa, A-239
shutdown, A-271
kconfig
Concentrix command, 5-2, 5-5
kennedy tape drive, 3-28
kennedy9100 tape drive, 3-28
kernel debugger
Concentrix-2800, 1-21
kernel
build custom, B-36
execution profile, A-110
maxsharesize, A-256
memory, 3-6
patch, 5-31, B-60
print version information, A-320
profiling, A-108
setclassdef, A-248
statistics server, 7-7, A-223
variable, B-60
keyenvoy
administration command, A-106
keylogin
Concentrix command, 2-3
keyserv
administration command, 7-2, A-107
keyserver
daemon, A-107
talk to, A-106
keystone tape drive, 3-28
kgmon
administration command, A-108
kgprof
administration command, A-110

kill
cluster, 5-19
Concentrix command, 1-39, 5-2, 5-8, 7-3
process, 5-8
klog
log device, 3-6
kmem
file, 5-1
kernel memory, 3-6

## L

L
diskf, A-31
rmt, A-210
l
diskf, A-34
-L
uucico, A-312
-l
fastboot, A-56
fasthalt, A-58
halt, A-88
implog, A-95
lpd, A-116
patchunix, A-179
reboot, A-193
sa, A-239
spray, A-277
L-devices file, 7-1, 7-18
L-dialcodes file, 7-1, 7-18
L.aliases file, 7-1, 7-18
L.cmds file, 7-1, 7-18
L.sys file, 7-1, 7-18
label
disk, 3-13, A-36, B-42, B-51
SCSI disk, 3-22
Label_disk
administration procedure, 3-2, 4-2, B-51
lastlog file, 1-10
layered product
install, A-274
report, A-196
layout
disk, 3-17
file system, 4-23
level
dump, 4-7, A-41
lib directory, 1-9
/lib directory, 1-9
line printer
control program, 6-1, 6-5, A-114
daemon, A-116
list
directory contents, 4-3
file, 5-10
file/directory, 1-12
special file, 3-5
load tape
Cipher, 3-30
Telex, 3-29
load
8mm cartridge, 3-25
Diagnostix-2800, 1-16
Exabyte, 3-25
local
area network, 7-5
directory, 1-9
lock
network daemon, A-113
lockd
administration command, A-113
log
device for syslogd, 3-6
IMP interpreter, A-95
system messages, A-292
LOGFILE file, 1-10
logfile
cron, A-27
logical name
create table, A-226
daemon, A-224
delete table, A-230
dump table, A-225
enter in table, A-230
look up in table, A-232
purge from table, A-233
remove from table, A-234
logical unit number, 3-5
logical
batch device, 6-16
block number, A-62
.login file, 2-6
login
accounting, A-2
activate for terminal, 3-42
Concentrix command, 2-3, 2-3
getty, A-86
Internet service, 7-6
name, B-15
remote server, A-206
remote, 7-6
shell, B-15
system, 2-3
terminal on/off, B-10
logout Concentrix command, 2-3
.logout file, 2-6
logout
Concentrix command, 2-3
look up
logical name in table, A-232
lookup
program, 7-6
lost+found directory
fsck, A-137
/lost+found directory, 1-9
low priority
scheduling, 5-27
lp0 device, 6-3
lpc
administration command, 6-1, 6-5, A-114
error messages, 6-8
lpd
administration command, 6-1, A-116
file, 6-1
queue, 6-3
lpd-acct file, 6-15
lpd-errs file, 1-10
lpf output filter, 6-14
lpf
file, 6-1
lpq
Concentrix command, 6-2, 6-4, 6-5
error messages, 6-8
lpr
Concentrix command, 6-2, 6-4
error messages, 6-9
lprm
Concentrix command, 6-2, 6-4, 6-5
error messages, 6-9
ls
Concentrix command, 1-12, 4-2, 4-3, 5-10
restore, A-200

## M

M
diskf, A-31
m
restore, A-199
-M
catman, A-16
sendmail, A-242
timed, A-297
-m
makedbm, A-121
makefs, A-124
pac, A-178
sa, A-239
showsched, A-270
syslogd, A-292
tunefs, A-306
uuclean, A-315
ypinit, A-330
yppasswd, A-333
ypwhich, A-342
machine number, 1-4
magtape remote protocol module, A-210
mail server
open connection, A-133
mail
alias, 2-8
arrival notification, 7-6
Concentrix command, 2-3, 2-7
encrypted, 2-8
facility, 2-7
file, 2-2
install on new system, B-49
print queue, A-119
queue, 2-7

rebuild data base
aliases file, A-165
send over Internet, 7-7, A-241
sendmail, 2-9
spooling directory, 2-7
start-up, 2-8
Mail.rc file, 2-1, 2-8
mailbox, 2-7
mailq
administration command, 2-2, A-119
.mailrc file, 2-1, 2-6, 2-8
mailstats
administration command, 7-2, A-120
maintain system files, 1-11
maintenance processor
VMEDP, 1-5
maintenance
disk, D-1
major numbers
pseudo terminal, 7-9
major
device number, 3-2, A-142
make
directory and parents, A-123
encryption key, A-128
file system, 4-4, A-124, B-58
hashed host table, A-136
hashed password table, A-143
hosts directory, A-121
NIS dbm, A-121
special file for pseudo terminal, A-146
special file, 3-4, A-142
striped partition, A-147
swapping partition, A-148
whatis data base, A-132
makedbm
administration command, 7-2, A-121
makefs
administration command, 4-1, 4-5, 4-21, 4-23, A-124, B-58
devices, 3-1
MAKEHOSTS
administration command, A-121
makekey
administration command, A-128
makeramdisk
administration command, 3-1, A-129
makewhatis
administration command, A-132
make_dirs
administration command, A-123
Make_file_system
administration procedure, 4-2, B-58
man
create output files, A-16
directory, 1-9
make data base, A-132
manipulate
routing table, A-212
map
automount
NFS, A-7
DARPA port to RPC program number, A-183
NIS
get version, A-334
transfer, A-343
mappedisable kernel variable, B-62
mapping
DARPA port to RPC program number, 7-6
mark partition, 3-18
mask
permissions, 2-6
master IOM, 1-5
master
NIS server, 7-19, B-20
NIS
determine, A-342
maxsharesize
kernel variable, B-63
kernel, A-256
MAX_SHARED_CREATE
kernel variable, B-63
MAX_SHARED_CREATE
setmaxshare, A-256
MAX_SHARED_PERMCRE ATE kernel variable, B-63
MAX_SHARED_PERMCRE ATE
setmaxshare, A-256
mbufs kernel variable, 7-16
mconnect
administration command, 7-2, A-133
MDQS, 6-16
me
timezone, A-301
media
dump, 4-10
mem
main memory, 3-6
memory map
Concentrix-2800, 3-44
memory module
hardware identification, 1-7
memory subsystem, 1-7
memory
configurations, 1-7
create shared region, A-281
delete shared region, A-282
Diagnostix-2800 memory map, 3-48
display shared region, A-283
kernel, 3-6
main, 3-6
restrictions, A-196
top users, 5-5
VMEbus memory map, 3-44
message
broadcast, B-34
log system, 7-7
of the day, B-34
system, 2-7, B-34
message-of-the-day, 2-7
messages file, 1-10, 1-23, 1-41, 5-1
messages
file, 1-1
metric
ifconfig, A-91
microcode type and revision level, A-196
migrate to IP
implicit, 5-29
mine
timezone, A-301
minimum free space
file system, 4-23
Minor number
rules for specifying, 3-34
minor
device number, A-142
number, 3-2
minw
timezone, A-301
mkdir
Concentrix command, 2-3, 4-2, 4-5
mkdisknod
administration command, 3-1, 3-4, A-134
Diagnostix-2800 devices, 3-10
mkfile
administration command, A-135
mkhosts
administration command, A-136
mklost+found
administration command, 4-1, A-137
mknewdevs script, 3-10, 7-9
mknewdevs
administration command, A-138
default devices, 3-36
mknod
administration command, 3-1, 3-4, 3-34, 3-36, 4-1, 7-2, 7-9, A-142
mkpasswd
administration command, 2-2, 2-4, A-143

mkproto
  administration command, A-144
mkptynod
  administration command, 7-9, A-146
mkstripe
  administration command, 3-1, 4-1, A-147
  file, 1-1
mkswap
  administration command, A-148
mktape
  administration command, 3-1, 3-4, 3-28, 3-34, A-149
mkttynod
  administration command, 3-1, 3-4, 3-36, A-150
mlp
  VMEbus printer, 3-6
mlp0 device, 6-3
MM
  configuration, 1-7
/mnt
  directory, 1-9
Model 200 PM
  configuration, 1-4
  scheduling, 5-11
Model 400 PM
  configuration, 1-4
  scheduling, 5-11
modem devices
  changing mknewdevs script, 3-36
modem
  connection, 3-39
  duart port, 3-6
modes of operation, 1-13
modify structure
  special tape drive, 3-32
modify
  cluster configuration, 5-20
  scheduling vector, 5-22
mon
  Concentrix command, 5-2, 5-4
monitor
  diagnostic, 1-18
  network status, A-284
  system, 5-1, 5-4
monthly script, 2-11, 5-30
motd, B-34
  Concentrix command, 2-7
  file, 1-1, 1-11, 2-1
mount
  administration command, 3-1, 4-1, 4-3, 4-5, 7-2, A-151
  file system, A-151, B-58
  Internet service, 7-6
  NFS filesystem automount, A-7
  NFS mount request server, A-155
  remote, 7-6
  root disk
    Concentrix, D-10
    Diagnostix, D-10
  show remote, A-269
mountd
  administration command, 7-2, A-155
move
  directory with bfio, 4-14
  directory with tar, 4-15
  file system with dump, 4-13
mqueue file, 2-2
msg directory, 1-9
.msgrc file, 2-1
msgs
  command, B-34
  Concentrix command, 2-3, 2-7
  file, 2-2
mt
  timezone, A-301
mtab
  file, 3-1, 4-1
  mount command, A-151
multi-user mode
  Concentrix-2800, 1-22
multi-user
  Concentrix, 1-8
  mode, B-23
    halt from, 1-14
    reboot from, 1-14
multiple-disk swapping/paging area, A-55
multiplexor
  Add_terminal, B-10
  connection, 3-39
  switch, 3-40
  terminal, 3-35, A-150
  VMEbus, 3-6, 3-38

## N

n
  diskf, A-34
  dump, A-42
-n
  catman, A-16
  fastboot, A-56
  fasthalt, A-58
  fsck, A-61
  halt, A-88
  makefs, A-124
  mkfile, A-135
  mount, A-151
  netdae, A-162
  nfsstat, A-169
  quot, A-185
  reboot, A-193, A-193
  route, A-212
  rpcinfo, A-217
  sa, A-239
  sendmail, A-242
  shutdown, A-271
  timed, A-297
  uuclean, A-315
name
  device, 3-4
  disk, 3-11
  domain, 7-19
  Ethernet device, 7-8
  network, 7-10
  remote user, 7-18
  server, A-156
  tape, 3-27
  terminal, 3-35
  uucp port, 7-18
named pipe
  mknod, A-142
named
  administration command, A-156
  Internet service, 7-6
names
  SCSI disk devices, 3-21
nbuf kernel variable, B-63
ND
  boot block server, A-161
ndbootd
  administration command, A-161
netdae
  administration command, 6-1, 6-17, A-162
netgroup
  ypmake, A-331, B-20
netmask
  ifconfig, A-91
.netrc file, 7-1
netrecv
  administration command, 6-1, 6-17, A-163
netsend
  administration command, A-164
netstat
  Concentrix command, 7-3, 7-8, 7-14
network information service, 7-19
  add to network, B-20
  data base values, 7-19
  files, 7-1
  setup, 7-20
network
  address, 7-10
  Add_ethernet, B-2
  batch receive daemon, A-163
  batch send daemon, A-164
  batch server daemon, A-162
  configuration, 7-11
  debugging, A-182
  dump, A-192
  examine and test, 7-14
  files and commands, 7-1
  hostname, 7-10
  interface parameter, A-90
  local area, 7-5

lock daemon, A-113
name, 7-10
restore, A-216
routing daemon, 7-7, A-214
rwall server, A-236
status monitor, A-284
status, 7-8
username server, 7-7, A-235
uucp, 7-17
networks
file, 1-11, 7-1, 7-11
ypmake, A-331, B-20
new clusters
scheduling vector, 5-25
newaliases
Add_user, B-13
administration command, 2-2, A-165, B-49, B-70
newkey
administration command, 7-2, A-166
nf
print filter, 6-14
nfile kernel variable, B-63
NFS
create swap file, A-135
daemon, A-14, A-168
export client, A-52
install diskless, A-102
mount request server, A-155
secure database, A-166
statistics, A-169
nfsd
administration command, A-168
nfsstat
administration command, A-169
nfs_wakeup_one_nfsd kernel variable, B-63
NIC host table, A-85
standard format, A-89
ninode kernel variable, B-63
NIS
add to network, B-20
binder process, 7-19, A-326
build and install, A-330
description, 7-19
determine master, A-342
files, 7-1
force map propagation, A-336
get map from master, A-335
get map version, A-334
make dbm, A-121
maps, 7-20
password file, A-333
point to particular server, A-340
rebuild, A-331, A-331
server process, A-338
server to change info, A-341
transfer map, A-343
update data base, B-22
nmbclusters kernel variable, B-63
nmbufs kernel variable, B-63
noauto
mount, A-151
nodst
timezone, A-301
nologin, A-271
file, 1-1, 2-1
non-interactive
file system check, A-62
non-network
set localhost, 7-12
NOPUSH
ypmake, A-331
nosuid
mount, A-151
notify
dump, A-42
npd
NFS remote execution driver, 3-6
nproc kernel variable, B-63
nptys kernel variable, A-184, B-63
nroff man pages, A-16
nrschs
8mm cartridge
no rewind, 3-6
default devices, 3-24
nrxtv
VMEbus tape controller, 3-6, 3-27
nslookup
administration command, A-171
ntalk
Internet service, 7-6
ntalkd
administration command, 7-2, A-176
null device, 3-6


## O


O
diskf, A-31
rmt, A-210
-o
dlabel, A-36
export client, A-52
makedbm, A-121
mount, A-151
sendmail, A-242
tunefs, A-306
-ob
dlabel, A-36
object file, 5-10
of
print filter, 6-14
off
terminal for login, B-10
on
terminal for login, B-10
on-line documentation, A-16
operations
basic, 1-1
operators
Label_disk, B-53
optimize
file system, A-306
options command, 1-13
console mode, 1-15
options
Concentrix-2800 mode, 1-22
ddt mode, 1-17
FX/2800 booter mode, 1-20
FX/2800 startup mode, 1-18
single-user Co-ncentrix-2800, 1-21
VMEbus booter mode, 1-16
opt_software
administration command, A-177
output
print filter, 6-14
ownership
change file owner, A-18
file system, A-185


## P


P
diskf, A-33
p
diskf, A-31, A-34
mknod, A-142
-P
pac, A-178
-p
ac, A-2
catman, A-16
edquota, A-50
fsck, A-61
kgmod, A-108
mount, A-153
named, A-156
netdae, A-162
pac, A-178
quotacheck, A-186
rpcinfo, A-217
trpt, A-302
trsp, A-304
uuclean, A-315
pac
administration command, A-178
packet
spray, A-277
page/swap
space problems, B-66

pagesharedflag kernel variable, B-63
paging
  configuration, A-196
  enable/disable, A-323
  extend area, B-40
  multiple-disk paging area, A-55
  partition, 3-14, 3-17, 3-18
pallocsize kernel variable, B-64
panel
  front, 1-3
panic
  recover, B-69
  string, 1-40
  system, 1-40
parallel
  printer, 6-3, 6-11, B-6
parameters command
  console mode, 1-15
PARC universal packet protocol, 7-5
partition
  add, B-55
  boot, 3-14
  default
    file system, 4-4
  disk, 3-5, 3-13
  make striped, A-147
  non-standard, 3-19
  remove unused, 3-18
  removing unused, 3-18
  SCSI disk, B-54
  swapping/paging, 3-14, 3-17, 3-18
  table fields, B-52
  unmarking, 3-18
passwd -f
  Add_user, B-13
  Concentrix command, 2-3
passwd
  chown, A-18
  Concentrix command, 2-3
  file, 1-1, 1-11, 2-1, 2-4, 4-1, 7-1
  format of file, B-15
  ypmake, A-331, B-20
passwd.dir
  file, 2-1, 2-4
passwd.pag
  file, 2-1, 2-4
password
  edit file, A-322
  make hashed table, A-143
  NIS, A-333
patch
  Concentrix image, A-179
  kernel, 5-31
patchunix
  administration command, A-179
Patch_kernel
  administration procedure, 5-1, B-60
pathname
  directory, 4-3
  remote system permissions, 7-18
performance
  Ethernet, 7-15
peripheral error, 1-41
permissions
  logical name table, A-226
pertec tape drive, 3-28
pg
  Interphase Peregrine FDDI, 3-6
phase 1
  fsck, A-69
phase 1b
  fsck, A-71
phase 2
  fsck, A-71
phase 3
  fsck, A-74
phase 4
  fsck, A-75
phase 5
  fsck, A-77
phase 6
  fsck, A-78
phase 7
  fsck, A-79
phone number
  remote system, 7-18
pin
  terminal connections, 3-38
ping
  administration command, 7-2, 7-8, 7-14, A-182
placement
  Ethernet driver, 7-15
.plan file, 2-1, 2-6
platter
  disk, 3-12
PLIST
  SCSI disk, 3-23
plot filter, 6-14
plotter
  accounting, A-178
PM
  hardware identification, 1-4
  software identification, 1-4
poll
  remote uucp site, 7-17, A-317
port
  mount, A-152
  terminal, 3-37
  uucp port name, 7-18
portmap
  administration command, 7-2, A-183
  Internet service, 7-6
POSIX
  timezone, A-345
power
  front panel, 1-3
  system, 1-14, B-25
power-up, 1-22
pp
  diskf, A-33
pr
  Concentrix command, 6-2
preserve directory, 1-9
print
  abort job, 6-6
  cluster information, 5-20, A-21
  Concentrix command, 6-2, 6-4
  control program, 6-5
  error messages, 6-7
  Ethernet board statistics, A-93
  examine queue, 6-4
  kernel version information, A-320
  mail queue, A-119
  NIS master info, A-342
  queue, 6-1, 6-3
  queueing
    disable, 6-6
    enable, 6-6
  rank job, 6-7
  remove job, 6-4, 6-7
  restart job, 6-7
  sendmail statistics, A-120
  start printing, 6-6
  stop printing, 6-6
  submit job, 6-4
  terminfo
    readable, A-49
  trouble-shooting, 6-7
printcap
  Add_printer, B-7
  file, 1-11, 6-1, 6-10
  lpd, A-116
printer file, 6-1
printer system, 6-3
PRINTER
  environment variable, 6-4
printer
  accounting, 6-15, A-178
  add to system, B-6
  control program, A-114
  default, 6-10
  output filter, 6-14
  parallel, 6-11
  remote, 6-13
  serial, 6-12
  user capabilities, 6-4
  VMEbus, 3-6
priority
  batch queue, 6-16
  high/low scheduling, 5-27
  process, A-195
problem
  report to Alliant customer service, B-37
procedures
  RAM disk utilities, D-4
process

change computing resource type, A-263
class, 5-13
control, A-100
kill, 5-8
priority, A-195
set class, A-246
status, 5-4, 5-6
thread, 5-6
processor module, 1-4
processor
attach, 1-30
default configurations, 5-14
display activity, 5-7
time, 2-10
product
install layered, A-274
profile
kernel, A-108, A-110
program
remote execution, 7-7
.project file, 2-1, 2-6
prompt
dunix, 1-17
FX/2800 booter mode, 1-20
FX/2800 startup mode, 1-18
VMEbus booter mode, 1-16
protection
disk, 3-18
protocol driver
Ultranet, 3-6
protocols
file, 7-1
ypmake, A-331, B-20
prototype
file system, A-144
ps
Concentrix command, 1-39, 5-2, 5-6
pseudo terminal, 3-6, 7-9
default devices, 7-9
major/minor numbers, 7-9
make special file, A-146
making device, 3-5
pty, 3-6
ptynum, A-184
tty, 3-6
pt
timezone, A-301
ptmp file, 2-1
ptmp
file, 2-4
pty
pseudo terminal, 3-6
client, 3-6
special file, 7-4
ptynum
administration command, 7-9, A-184
pub directory, 1-9
publickey
create new key, A-166
file, 7-1
pull NIS map, A-335
PUP
PARC universal packet protocol, 7-5
purge
logical name table, A-233
push NIS map, A-336
pwd
restore, A-200

## Q

-q
fastboot, A-56
fasthalt, A-58
halt, A-88
reboot, A-193
routed, A-214
sendmail, A-242
qconf file, 6-1, 6-17, B-46
batchd, A-12
qconf
file, 1-11
qdel
Concentrix command, 6-2, 6-18
qstat
Concentrix command, 6-2, 6-18
qsub
Concentrix command, 6-2, 6-17
query
Internet name servers, A-171
queue
batch, 6-1, 6-16, A-12
examine print, 6-4
print mail, A-119
print, 6-1, 6-3
quit
lpc, A-115
restore, A-200
timedc, A-299
quot
administration command, 4-1, A-185
quota
Concentrix command, 2-3, 4-2
data base, 2-13
filesystem
turn on/off, A-187
remote server, A-215
summary for file system, A-197
quota/noquota
mount, A-152
quotacheck
administration command, 2-2, 4-1, A-186
quotaon
administration command, 2-2, 4-1, A-187
quotas
boot script, 1-33
file sysem
add to system, B-8
file, 2-1, 2-13, 4-1
user
edit, A-50

## R

R
restore, A-198
rmt, A-210
r
diskf, A-33
restore, A-198
-R
chown, A-18
sendmail, A-242
uucico, A-312
-r
ifstats, A-93
implog, A-95
kgmod, A-108
mount, A-153
nfsstat, A-169
pac, A-178
ping, A-182
sa, A-239
sendmail, A-242
shutdown, A-271
uucico, A-312
RAM disk image
create, A-129
RAM disk
block device, A-129
boot from, 1-37
devices, D-9
Diagnostix
install utilities, D-3
maintenance utilities, D-2
raw device, A-129
srboot, A-275
utilities, D-1
procedures, D-4
ramdisk
make, A-129
srboot, A-275
random
i-node generation numbers, A-81
rank
print job, 6-7
Raster filter, 6-14
raw device mode, A-142
raw device
RAM disk, A-129
rc
administration command, A-189
script, 1-1, 1-2, 1-23, 1-32
rc.diag script, 1-1, 1-31, A-100, A-189

rc.local script, 1-1, 1-11, 1-23, 1-34, 7-1, A-189
rcp
  Concentrix command, 7-3
rdate
  adminisratation command, 7-2
  administration command, A-191
rdump
  administration command, 7-2, 7-8, A-192
retime scheduler, 5-29, A-259
realtime protection device, 3-6
realtime retime protection device, 3-6
reboot, A-271, 1-2, A-193
  Boot_system, B-23
  from multi-user mode, 1-22
  multi-user mode, 1-14
  system, A-56, A-56, A-193
rebuild
  mail aliases data base, A-165
reconfiguring
  cluster
    dynamically, 5-20
recover
  page/swap problems, B-66
  system, 1-39
Recover_system
  administration procedure, 1-2, B-66
relabel disk, 3-13, B-51
relabel
  disk stripe, B-57
remote execution
  server, A-204
remote host
  set system date, A-191
remote login
  Apollo window, 3-43
remote system
  add, B-18
remote user
  information server, A-59
remote
  batch queue, 6-16
  dump, A-192
  execution, 7-6
  file calling information, 7-18
  file, 1-11, 7-1
  login, 7-6
    server, A-206
  magtape protocol module, A-210
  mail from uucp handle, A-209
  mount file system, 7-6
  pathname permissions, 7-18
  poll uucp site, A-317
  print job, 6-3
  printer, 6-13
  program execution, 7-7
  quota server, A-215
  restore file system, A-216
  RPC execution server, A-203
  shell server, A-221
  shell, 7-7
  show mount, A-269
  user name, 7-18
  uucp site, 7-17
remount
  mount, A-152
remove
  cluster, 5-20, A-23
  diskless client, A-264
  logical name table, A-234
  print job, 6-4, 6-7
  unused partitions, 3-18
  user, B-70
Remove_user administration procedure, 2-2
Remove_user
  administration procedure, B-70
renice
  administration command, A-195
repair
  file system, A-61
replicated superblock
  file system, 4-26
report
  RPC information, A-217
  software/hardware configuration, A-196
  system accounting, A-239
reportcfg
  administration command, A-196
repquota
  administration command, 2-14, 4-1, A-197
reset
  front panel, 1-3
  hung system, 1-22
  system, 1-14, B-25
resource
  accounting, 2-10
  dedicate, 5-28
  set scheduling vector, A-259
restart
  from panic, B-69
  lpc, A-115
  print job, 6-7
  system, A-193
restore
  administration command, 4-1, 4-7, A-198
  Concentrix-2800 root partition, 4-18
  diagnostics, A-201
  dump across network, 7-8
  file system, 4-7, 4-11
  incremental file system, A-198
  interactive, 4-12
  remote file system, A-216
  root partition, 4-16
  selectively, 4-12
retrans
  mount, A-152
retry
  mount, A-152
rexd
  administration command, A-203
  Internet service, 7-7
rexecd
  administration command, 7-2, A-204
  diagnostics, A-205
rf
  print filter, 6-14
rh
  diskf, A-33
.rhosts file, 7-1
ri
  diskf, A-33
rlimit_core kernel variable, B-64
rlogin Concentrix command, 2-3
rlogin
  Concentrix command, 2-3, 7-3
rlogind
  administration command, 7-2, A-206
  control terminal, A-206
  diagnostics, A-208
rmail
  administration command, A-209
  Concentrix command, 7-3
rmdk
  physical memory disk, 3-6
rmt
  administration command, 7-2, A-210
rmtab
  file, 7-1
ro
  export client, A-52
  mount, A-151
root disk
  mount, D-10
root file system
  RAM disk, 1-37
root partition
  Concentrix-2800, 1-25
  diagnostix-2800, 1-24
  restoring, 4-16
root, 1-8
  create for diskless, A-26
  directory, 1-9
  diskless, A-102
  export client, A-52
  file system, 3-17, 4-3

rotational delay
  file system, 4-23, 4-26
  tunefs, A-306
round-robin
  batch queue, 6-16
route
  administration command, 7-2, A-212
  Internet service, 7-7
routed
  administration command, 7-2, A-214
routing
  network, 7-7, A-214
  table manipulation, A-212
RPC
  information, 7-8
  keyenvoy command, A-106
  map number from DARPA port, A-183
  remote execution server, A-203
  remote quota server, A-215
  report information, A-217
  statistics, A-169
rpc
  ypmake, A-331, B-20
rpcinfo
  administration command, 7-8, A-217
rquotad
  administration command, A-215
rrestore
  administration command, 7-2, 7-8, A-216
rschs
  8mm cartridge, 3-6
  default devices, 3-24
rsd
  SCSI disk device, 3-6
  SCSI disk drive defaults, 3-22
rsh
  Concentrix command, 7-3
rshd
  administration command, 7-2, A-221
  diagnostics, A-222
rsize
  mount, A-152
rstatd
  administration command, A-223
  Internet service, 7-7
rstripe striped device, 3-6
rt
  Concentrix command, 3-31
rtnamed
  administration command, A-224
rtncrtbl
  administration command, A-226
rtndeltbl
  administration command, A-230
rtndumptbl
  administration command, A-225
rtnenter
  administration command, A-230
rtnlookup
  administration command, A-232
rtnpurge
  administration command, A-233
rtnremove
  administration command, A-234
RTS function
  cabling, 3-38
rules
  minor number for terminal, 3-37
ruptime
  Concentrix command, 7-3
rusersd
  administration command, A-235
  Internet service, 7-7
rw
  export client, A-52
  mount, A-151
rwall
  Concentrix command, 7-3
rwalld
  administration command, 7-2, A-236
rwho
  Concentrix command, 7-3
rwhod
  administration command, 7-2, A-237
RXD function
  cabling, 3-38
rxdv
  VMEbus disk, 3-6, 3-11
rxtv
  VMEbus tape, 3-6, 3-27

## S

S
  diskf, A-33
  rmt, A-210
s
  diskf, A-33, A-34
  dump, A-42
  restore, A-199
-S
  fastboot, A-56
  makefs, A-124
  reboot, A-193
  sa, A-239
  shutdown, A-271
-s
  arp, A-5
  fastboot, A-56
  kgprof, A-110
  nfsstat, A-169
  pac, A-178
  reboot, A-196
  routed, A-214
  sa, A-239
  shutdown, A-271
  trpt, A-302
  trsp, A-304
  uucico, A-312
  ypinit, A-330
  ypupdated, A-341
Sa
  diskf, A-33
sa
  administration command, 2-2, 2-10, A-239
safety measure
  diskf, 3-18
savacct, A-239
  file, 1-10, 2-1, 2-10
save directory
  bfio, 4-14
  tar, 4-15
scheduler
  retime, 5-29, A-259
  time-sharing, A-259
scheduling vector
  clcr, 5-25
scheduling
  computational resources, 5-11, 5-13
  display default classes, A-267
  display vector, A-270
  modify vector, 5-22
  program execution, 5-12
  set default class, A-254
  set resource, A-259
  vector, 5-14
script
  boot, 1-25
  crontab, 5-30
  for dumps, 4-9
  rc, 1-1, 1-2, 1-23, 1-32
  rc.diag, 1-1, 1-31
  rc.local, 1-1, 1-11, 1-23, 1-34
  start system, A-189
SCSI adaptor number
  Diagnostix-2800, 3-10
SCSI bus
  making device, 3-5
SCSI disk device, 3-6
SCSI disk drive interface, B-42
SCSI disk
  Label_disk, B-54
SCSI minor number
  Diagnostix-2800 devices, 3-10
SCSI

8mm cartridge names, 3-21
devices, 3-21
disk labels, 3-22
drives supported, 3-21
flaws, 3-23
supported disk drives, 3-22
switch configuration, 3-24
tape devices, 3-24
sd
SCSI disk device, 3-6
secretmail
directory, 2-8
file, 2-2
sector
disk, 3-12
sectors per track, A-124
slip, 3-16
secure terminal, 3-42, B-10
secure
export client, A-53
mount
NFS, A-152
RPC
upupdated, A-341
security
NFS filesystem, A-81
select
diskless software from distribution, A-177
selective file system restore, 4-12
send
icmp echo packet, A-182
mail over Internet, 7-7, A-241
sendmail
administration command, 2-2, 7-2, A-241, B-49
configuration options, A-242, C-45
diagnostics, A-244
facility, 2-9
implementation, C-6
installation, C-15
operation, C-4, C-15
overview, C-3
tuning, C-27
usage, C-6
uucp network, 7-17
sendmail.cf file, 1-11, 2-1, 2-2
sendmail.hf file, 2-1
sendmail.main.cf file, 2-9
sendmail.st file, 2-1
sendmail.subsidiary.cf file, 2-9
senmail
print statistics, A-120
serial
printer, 6-3, 6-12, B-6
server
DARPA
Internet file transfer protocol, A-82
diskless install, A-102
network information service, A-338
NIS, 7-19, B-20
servers
file, 7-1
services
file, 1-11, 7-1
ftpd, A-82
ypmake, A-331, B-20
set
computing resource, A-252
control terminal
rlogind, A-206
default scheduling class, A-254
disk quotas, 2-13
localhost
non-network machine, 7-12
maximum no. of processors concurrent job, A-258
NIS, A-340
process class, A-246
scheduling vector for computing resource, A-259
system date
from remote host, A-191
terminal mode, A-86
threshold for computation complex, A-250
time zone, A-301
timezone, A-345
setclass command
scheduling class, 5-13
setclass system call
scheduling class, 5-13
setclass
administration command, 5-1, 5-13, A-246
setclassdef kernel variable, B-64
setclassdef
administration command, 5-1, A-248
setcomplex
administration command, 5-17, A-250
setcres system call
scheduling class, 5-13
setcres
administration command, A-252
system call, 5-12
setdefclass
administration command, 5-13, A-254
setmaxshare
administration command, A-256
setmodes
restore, A-200
setpn
administration command, A-258
setsched
administration command, 5-1, 5-22, A-259
settype
administration command, A-263
setup
batch, 6-17
network information service, 7-20
setup_client
administration command, A-264
setup_exec
administration command, A-266
shareable executable file
sticky, A-285
shared
create memory region, A-281
delete memory region, A-282
display memory region, A-283
maximum memory, A-256
region
limits, A-256
shell
Internet service, 7-7
login, B-15
remote server, A-221
remote, 7-7
show
remote mounts, A-269
showdefclass
administration command, A-267
showmount
administration command, 7-2, A-269
showsched
administration command, 5-1, 5-15, 5-22, A-270
shutdown
administration command, 1-2, A-271, B-45
Boot_system, B-23
system to single-user mode, B-73
system, 1-14, A-271
to diagnostix-2800 mode, 1-17
to Diagnostix-2800 mode, 1-22
to single-user mode, 1-22, B-24
Shutdown_system
administration procedure, 1-2, B-73
signal ground, 3-38
single-user mode
Boot_system, B-32
Concentrix-2800, 1-21
single-user
Concentrix, 1-8

fastboot, A-56
mode, 1-14
Boot_system, B-24
shutdown, B-73
reboot, A-196
size
dump, 4-10, A-42
partition in sectors, A-124
standard disk partitions, 3-14
skel
directory, 1-9, 2-6
files, 2-2
skeleton files, B-13
skeleton
file, 2-6
slave
NIS server, 7-19, B-21
slice
time, 5-14
slip bad sector
SCSI disk, 3-23
slip sector, 3-16, B-42
SMD-E disk drive interface, B-42
smtp Internet service, 7-7
snapshot
uucp system, 7-17, A-318
soft
mount, A-152
software
configuration, 1-8
terminal, 3-42
network problem, 7-14
report configuration, A-196
software-install
administration command, A-274
soft_net_cpu kernel variable, B-64
source
Concentrix files, B-36
special file
arp, 7-4
build, A-142
default, 3-10
description, 3-2
icmp, 7-4
if, 7-4
inet, 7-4
ip, 7-4
list, 3-5
make for disk, A-134
make for pseudo terminal, A-146
make for tape, A-149
make for terminal, A-150
make, 3-4
pty, 7-4
tcp, 7-4
tty, 7-4
udp, 7-4
special tape drive, 3-32
special tape drives
minor number, 3-34
specifications
configuration file, 5-19
spool
directory, 1-9, 6-1
Add_printer, B-7
clean up, 7-17
print job, 6-3
spray
administration command, A-277
daemon, A-278
sprayd
administration command, A-278
Internet service, 7-7
-Sr
dlabel, A-36
srboot administration command, A-275
srbootflags
file for RAM disk, A-129
src
directory, 1-9
files, 3-1
srcreate
administration command, A-281
srdelete
administration command, A-282
srdisplay
administration command, A-283
-Ss
dlabel, A-36
stackbottom kernel variable, B-64
standard partitions
sizes, 3-14
start
lpc, A-115
mail, 2-8
printing, 6-6
Sun system, A-286
system, A-56, A-189, B-23
Startup and shutdown procedures, 1-22
startup file
cluster configuration, 5-21
startup
alliant ddt 2800 mode, 1-21
cluster configuration, 5-18
ddt mode, 1-17
Diagnostix-2800 mode, 1-17
FX/2800 booter mode, 1-20
FX/2800 startup mode, 1-18
modes of operation, 1-13
multi-user mode, 1-22
single-user Concentrix-2800, 1-21
VMEbus booter mode, 1-16
statd
administration command, A-284
statically detach CE, 5-13
statistics
ethernet board, A-93
kernel, 7-7, A-223
NFS, A-169
RPC, A-169
xdv disk, A-325
status
batch, 6-18
lpc, A-115
network, 7-8
process, 5-4, 5-6
system, 5-4, 7-7
server, A-237
stc tape drive, 3-28
sticky bit
description, A-285
sticky
administration command, 1-2
stop
lpc, A-115
printing, 6-6
storage layout
disks, 3-12
store
keys
server, A-107
string
panic, 1-40
stripe
disassemble disk, A-310
striped device, 3-6
striped device, 3-6
striped
file system, 4-4, B-57, B-58
make partition, A-147
stripeswapflag kernel variable, B-64
stripetab
file, 3-1
su
Concentrix command, 1-8, 2-3, 2-3
setclassdef, A-248
submit
batch request, 6-17
print job, 6-4
substitute user, 1-8, 2-3
subsystem
memory, 1-7
suid
mount, A-151
summary
file system information, 4-26
file system ownership, A-185
file system quotas, A-197
quotas, 2-14
Sun
boot system, A-286

sunboot
administration command, A-286
super-user, 2-3
Concentrix, 1-8
superblock
alternate
tunefs, A-306
file system, 4-22, 4-26
fsck, A-61, A-64
information, A-45
update, A-290, A-311
supported devices, 3-5
Sv
diskf, A-33
swap
create NFS swap file, A-135
diskless, A-102
swapping
configuration, A-196
enable/disable, A-323
extend area, B-40
make partition, A-148
multiple-disk swapping area, A-55
partition, 3-14, 3-17, 3-18
swdevt kernel variable, B-64
swdevt+0x4 kernel variable, B-64
switch configuration
SCSI, 3-24
switch
hardware identification, 1-7
multiplexor, 3-40
swpgratio kernel variable, B-64
sync
administration command, 4-1, 4-22, A-290
synchronization
file system, 4-22
sys5 directory, 1-9
syscfg
Concentrix command, 3-10, 5-2, 5-3
SYSLOG file, 1-10
syslog
administration command, 7-2
Internet service, 7-7
syslog.conf file, 2-1
syslog.pid file, 2-1
syslogd
administration command, 2-2, A-292
system
accounting, 2-10, A-3, A-239
active, 5-3
add new user, B-13, B-13
add printer, B-6
add terminal to, B-10
add to Ethernet, B-2
add to uucp network, B-17
configuration, 5-3
directories, 1-9
equivalent, 7-13
files that grow, 1-10, 5-10
files to maintain, 1-11
halt in console mode, A-58, B-45
hang, 1-40
hardware configuration, 5-11
initialization and process control, A-100
initialization, 1-31
log messages, A-292
login, 2-3
memory usage
display, 5-5
message, B-34
messages, 2-7
modes of operation, 1-13
monitor, 5-1
panic, 1-40
power off/on, B-25
reboot, A-56, A-193
report configuration, A-196
reset, B-25
shut down to single-user mode, B-73
shut down, 1-14, A-271
start, 1-14, B-23
start-up script, A-189
status server, A-237
status, 5-4, 7-7
tune, 5-1
system.config file, 1-11
sys_prefetch kernel variable, B-64

## T

T
diskf, A-33
t
diskf, A-33
restore, A-198
-T
makefs, A-124
-t
edquota, A-50
implog, A-95
lockd, A-113
makefs, A-124
mount, A-153
routed, A-214
rpcinfo, A-217
sa, A-239
sendmail, A-242
setsched administration command, 5-17
setsched, A-259
timed, A-297
trpt, A-302
trsp, A-304
umount, A-308
uucico, A-312
ypwhich, A-342
table
Diagnostix-2800 devices, 3-9
talk
Concentrix command, 7-3
daemon, A-176
Internet service, 7-7
to another user, 7-7
to keyserver, A-106
to other user, 7-6
tape
archiver, 4-15
configuration, 3-28
devices, 3-27
examine, 3-31
load Cipher, 3-30
load Telex, 3-29
make special file, A-149
name, 3-27
operations, 3-29
special drive, 3-32
unload Cipher, 3-30
unload Telex, 3-29
VMEbus, 3-6
tar
Concentrix command, 4-2, 4-15
-tc
setsched administration command, 5-17
setsched, A-259
TCP
Internet Protocol, 7-5, A-82
tcp
special file, 7-4
tcp_recvspace kernel variable, B-64
tcp_sendspace kernel variable, B-64
-td
setsched administration command, 5-17
setsched, A-259
telex125
tape drive, 3-27
telex50
tape drive, 3-27
TELNET
DARPA TELNET protocol server, A-294
protocol, 7-7
telnet
Concentrix command, 2-3, 2-3, 7-3
Internet service, 7-7
telnetd
administration command, 7-2, A-294
temporary file, 5-10
termcap file
terminals, 3-43
termcap
Add_terminal, B-10

terminal
activate login, 3-42
add to system, B-10
Apollo window, 3-43
characteristics, B-10
connection, 3-39
console, 1-3
devices, 3-35
disable, 3-43
enable, 3-43
enable/disable, B-39
hardware configuration, 3-37
make special file, A-150
making device, 3-5
multiplexor, 3-35
name, 3-35, 3-42
port, 3-37
pseudo, 3-6, 7-9
rules for minor number, 3-37
secure, 3-42
set mode, A-86
software configuration, 3-42
type, 3-42, B-10
window control, 3-42
terminfo
compile database, A-24
print in readable form, A-49
test
network, 7-14
Tex filter, 6-14
tf
print filter, 6-14
-tf
setsched, A-259
-tfm
setsched, A-259
tftp
Concentrix command, 7-3
Internet service, 7-7
tftpd
administration command, 7-2, A-295
thorn tape drive, 3-28
thread
process status, 5-6
threshold
file system, A-306
ticks
setsched, A-259
time limits
filesystem
edit, A-50
time
cron execution, A-27
daemon, A-297
execute at specified time, 5-30
processor, 2-10
server, 7-7
set zone, A-301
set, 1-38
slice, 5-14
time-sharing
scheduler, A-259
timed
administration command, A-297
control program, A-299
Internet service, 7-7
timedc
administration command, A-299
timeo
mount, A-152
timezone
administration command, A-301
compiler, A-345
tip command, 1-18
tip
Concentrix command, 7-3
srboot, A-275
tmp directory, 1-9
/tmp directory, 1-9, 4-5
/tmp
cleaning up, 4-19
tnamed
administration command, A-296
topq
lpc, A-115
-tr
setsched, A-259
trace
timedc, A-299
transliterate protocol, A-302
transliterate sequenced packet protocol, A-304
track
disk, 3-12
number per cylinder, A-124
track-allocated
file system, 4-4, A-124
trailers
ifconfig, A-91
-trailers
ifconfig, A-91
transceiver, 7-5
Add_ethernet, B-2
transfer
file, 7-6
NIS map, A-343
uucp files, A-312
transliterate
protocol trace, A-302, A-304
transmission control protocol, 7-5
trivial file transfer, 7-7
protocol server, A-295
trivial name server
DARPA, A-296
troff filter, 6-14
trouble-shooting
print, 6-7
trpt
administration command, A-302
trsp
administration command, A-304
trusted systems and accounts, 7-13
tty
process's terminal, 3-6
pseudo terminal, 3-6
special file, 1-1, 7-4
terminal, 3-35
VMEbus multiplexor, 3-6
ttyC duart port, 3-6
ttyR duart port, 3-6
ttys
Add_printer, B-6
file, 1-1, 1-11, 2-1, 3-1, 3-42, 6-1, 7-1
format of file, B-10
getty command, A-86
tune
file system, 4-21, 4-23, A-306
system, 5-1
tunefs
administration command, 3-1, 4-1, 4-21, 4-23, A-306
TXT function
cabling, 3-38
tz+4 kernel variable, B-65

## U

u
dump, A-42
restore, A-199
-U
sa, A-240
-u
export client, A-52
kgmod, A-108
makedbm, A-121
newkey, A-166
patchunix, A-179
rpcinfo, A-217
sa, A-240
ucb system directory, 1-9
UDP
Internet protocol, 7-5
udp
special file, 7-4
udp_recevspave kernel variable, B-65
udp_sendspace kernel variable, B-65
Ultranet
hardware driver, 3-6
internet driver, 3-6
protocol driver, 3-6
umask
Concentrix command, 2-6
umount
administration command, 4-1, 4-6, A-308

unet
Ultranet protocol driver, 3-6
unexport
NFS client, A-52
unit
disk, 3-12
tape device, 3-5
unix command
console mode, 1-15
/unix file, 1-1, 5-1
unix kernel file, B-36, B-60
unix to unix
command execution, 7-17
copy, 7-17
unix
srboot, A-275
/unix
booter, 1-20
unload tape
Cipher, 3-30
Telex, 3-29
unloading
8mm cartridge, 3-26
Exabyte, 3-26
unmark partition, 3-18
unmkstripe
administration command, 4-1, 4-6, A-310
unmount
file system, A-308
unpst_recspace kernel variable, B-65
unpst_sendspace kernel variable, B-65
unserviced job load
setcomplex, A-251
up
ifconfig, A-91
lpc, A-115
update
administration command, 4-1, 4-22, A-311
NIS data base, B-22
NIS information, A-341
superblock, A-290, A-311
us
timezone, A-301
usage
disk, 5-9
user
add to system, B-13
communications, 2-7
Concentrix, 1-8
current users on system, 5-3
disk space, 3-17
identification number, B-15
lookup program, 7-6
name in remote system, 7-18
remove, B-70
set quotas on file system, B-8
USERFILE
Add_uucp, B-17
file, 7-1, 7-18
usermem kernel variable, B-65
username
network server, 7-7, A-235
users
cpu, 5-4
/usr file system, 4-3
/usr/bin directory, 1-9
/usr/etc directory, 1-9
usracct file, 1-10, 2-1
usracct, A-239
usrclass
ypmake, A-331, B-20
usrstack kernel variable, B-65
usrstacksegs kernel variable, B-65
utilities
disk maintenance, D-1
utility
disk maintenance, D-1
utmp file, 1-1, 1-10, 2-1
uu.daily, B-18
file, 7-18
uu.hourly, B-18, 7-18
UUAIDS
Add_uucp, B-17
uucico
administration command, 7-2, 7-17, A-312
uuclean
administration command, 7-2, 7-17, A-315
uucp
add system to, B-17
administration command, 7-17
Concentrix command, 7-3
daemon, 7-17, A-316
directory spool clean up, 7-17
execution file interpreter, 7-17, A-319
network, 7-17
poll remote site, A-317
poll remote uucp site, 7-17
remote site, 7-17
snapshot of uucp system, 7-17
snapshot, A-318
spool directory clean up, A-315
transfer files, A-312
uucpd
administration command, 7-2, 7-17, A-316
uulog
Concentrix command, 7-3
uuname
Concentrix command, 7-3
uupoll
administration command, 7-2, 7-17, A-317
uuq
Concentrix command, 7-3
uusend
Concentrix command, 7-3
uusnap
administration command, 7-2, 7-17, A-318
uux
administration command, 7-17
commands executed by, 7-18
Concentrix command, 7-3
transfer files, A-312
uuxqt
administration command, 7-2, 7-17, A-319
uv
administration command, A-320
uvm
Ultranet hardware driver, 3-6

## V

v
diskf, A-33, A-34
restore, A-199
-V
dlabel, A-36
-v
catman, A-16
compile, A-24
export client, A-52
gettable, A-85
mkfile, A-135
mount, A-153
ndbootd, A-161
ping, A-182
quot, A-185
quotacheck, A-186
quotaon/off, A-187
repquota, A-197
sa, A-240
sendmail, A-242
umount, A-308
yppush, A-336
-V1
ypset, A-340
ypwhich, A-342
-V2
ypset, A-340
ypwhich, A-342
variable
kernel, 5-31, B-60
variables
Label_disk, B-53
ve
Interphase DMA exerciser, 3-6
vector
display scheduling, A-270
scheduling, 5-14
set scheduling, A-259
verbose

restore, A-200
verify
ecc on disk, A-30
mounted distribution tape, A-321
verify_tapevol_arch
administration command, A-321
versionname kernel variable, B-65
vf
print filter, 6-14
vipw
Add_user, B-13
administration command, 2-2, 2-4, A-322
VMEbus booter mode, 1-16
Boot_system, B-27
prompts, 1-16
VMEbus booter, 1-13, 1-14
VMEbus SCSI host adaptor, 3-21
VMEbus
cartridge, 3-6
disk, 3-6, B-42
memory map, 3-44
multiplexor, 3-6, 3-38
printer, 3-6
tape, 3-6
Xylogics tape controller, 3-27
VMEDP, 1-14
connect console, B-38
diagnostics subsystem, 1-17
maintenance processor, 1-5
vmlock
administration command, 5-1, A-323
/vmunix file, 1-1


## W


W
dump, A-42
rmt, A-210
w
Concentrix command, 5-2, 5-3
diskf, A-33
dump, A-42
-w
ac, A-2
catman, A-16
dlabel, A-36
kgprof, A-111
rtncrtbl, A-226
wall
command, B-34
Concentrix command, 2-3, 2-7
Internet service, 7-7
we
timezone, A-301
weekly script, 5-30
whatis
data base, A-16
make data base, A-132
who
Internet service, 7-7
wi
diskf, A-34
window
terminal control, 3-42
terminal, B-10
write
all users, 2-7
to all users, 7-7
to all, B-34
wsize
mount, A-152
wtmp
ac, A-2
file, 1-1, 1-10, 2-1, 2-11


## X


x
diskf, A-34
restore, A-198
-x
uucico, A-312
ypwhich, A-342
xdrtoc
administration command, A-324
xdv
VMEbus disk, 3-6, 3-11
xdvstat
administration command, A-325
xget
Concentrix command, 2-3
xl
diskf, A-34
xproc
command, 5-7
Concentrix command, 5-2
xsend
Concentrix command, 2-3
xtab
export client, A-52
mountd, A-155
Xylogics
754 disk controller, B-42
772 tape controller, 3-27
780 terminal multiplexor, 3-35
7800 disk controller, B-42
disk controllers, 3-11


## Y


y
restore, A-199
-y
fasthalt, A-58
fsck, A-61
halt, A-88
ypbind
administration command, 7-2, 7-19, A-326
ypcat
Concentrix command, 7-3, 7-19
ypinit
administration command, 7-2, A-330
ypmake
administration command, 7-2, A-331
yppasswd
administration command, 7-2, A-333
yppoll
administration command, 7-2, A-334
yppull
administration command, A-335
yppush
administration command, 7-2, A-336
ypserv
administration command, 7-2, 7-19, A-338
ypset
administration command, A-340
ypupdated
administration command, A-341
administration command, 7-3
ypupdaters
file, 7-1
ypwhich
administration command, 7-2, A-342
ypxfr
administration command, 7-3, A-343


## Z


z
diskf, A-34
-z
kgprof, A-111
nfsstat, A-169
zero i-node, A-22
zic
administration command, A-345
zone
set time, A-301

## READER'S COMMENTS

Please use this form to note suggestions for improvement, or any errors you find in this manual. Alliant Technical Publications Department will try to correct any errors in the next revision of the manual.

**Comments on** ***Concentrix-2800 System Administration - D002.18***

**Page** **Description**

Name _____ Date ___ _____

Company _____

Street _____

City _____ State _____ Zip _____

Telephone _____ Email _____

Outside USA, return free of charge to your local Alliant Customer Service office. You can also use email; send comments to docmgr@alliant.com.

Fold Here and Tape ----------

No postage
Necessasry
if mailed in the
United States

# BUSINESS REPLY MAIL

FIRST CLASS PERMIT # 118 LITTLETON MA 01460

POSTAGE WILL BE PAID BY ADDRESSEE

ALLIANT

ComputerSystemsCorporation

Technical Publications Dept.
One Monarch Drive
Littleton, Massachusetts 01460

-- Fold Here - Do not tear ---------- cut along dotted line this side >>>